

Simons Symposia

Matthew Baker
Sam Payne *Editors*

Nonarchimedean and Tropical Geometry

 Springer

Simons Symposia

More information about this series at <http://www.springer.com/series/15045>

Matthew Baker • Sam Payne
Editors

Nonarchimedean and Tropical Geometry

 Springer

Editors

Matthew Baker
School of Mathematics
Georgia Institute of Technology
Atlanta, GA, USA

Sam Payne
Department of Mathematics
Yale University
New Haven, CT, USA

ISSN 2365-9564

Simons Symposia

ISBN 978-3-319-30944-6

DOI 10.1007/978-3-319-30945-3

ISSN 2365-9572 (electronic)

ISBN 978-3-319-30945-3 (eBook)

Library of Congress Control Number: 2016942131

© Springer International Publishing Switzerland 2016

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper

This Springer imprint is published by Springer Nature
The registered company is Springer International Publishing AG Switzerland

Preface

1 Introduction

This volume grew out of two Simons Symposia on “Nonarchimedean and tropical geometry” which took place on the island of St. John in April 2013 and in Puerto Rico in February 2015. Each meeting gathered a small group of experts working near the interface between tropical geometry and nonarchimedean analytic spaces for a series of inspiring and provocative lectures on cutting edge research, interspersed with lively discussions and collaborative work in small groups. Although the participants were few in number, they brought widely ranging expertise, a high level of energy, and focused engagement. The articles collected here, which include high-level surveys as well as original research, give a fairly accurate portrait of the main themes running through the lectures and the mathematical discussions of these two symposia.

Both tropical geometry and nonarchimedean analytic geometry in the sense of Berkovich produce “nice” (e.g., Hausdorff, path connected, locally contractible) topological spaces associated to varieties over valued fields. These topological spaces are the main feature which distinguishes tropical geometry and Berkovich theory from other approaches to studying varieties over valued fields, such as rigid analytic geometry, the geometry of formal schemes, or Huber’s theory of adic spaces. All of these approaches are interrelated, however, and the papers in the present volume touch on all of them. The topological spaces produced by tropical geometry and Berkovich’s theory are also linked to one another; in many contexts, nonarchimedean analytic spaces are limits of tropical varieties, and tropical varieties are often best understood as finite polyhedral approximations to Berkovich spaces.

Topics of active research near the interface between tropical and nonarchimedean geometry include:

- Differential forms, currents, and solutions of differential equations on Berkovich spaces and their skeletons
- The homotopy types of nonarchimedean analytifications

- The existence of “faithful tropicalizations” which encode the topology and geometry of analytifications
- Relations between nonarchimedean analytic spaces and algebraic geometry, including logarithmic schemes, birational geometry, and linear series on algebraic curves
- Adic tropical varieties relate to Huber’s theory of adic spaces analogously to the way that usual tropical varieties relate to Berkovich spaces
- Relations between non-archimedean geometry and combinatorics, including deep and fascinating connections between matroid theory, tropical geometry, and Hodge theory

2 Contents

The survey paper of *Gubler* presents a streamlined version of the theory of *differential forms and currents on nonarchimedean analytic spaces* due to Antoine Chambert-Loir and Antoine Ducros, in the important special case of analytifications of algebraic varieties. Starting with the formalism of superforms due to Lagerberg, Gubler establishes or outlines the key results in the theory, including nonarchimedean analogs of Stokes’ formula, the projection formula, and the Poincaré–Lelong formula. Gubler also proves that these formulas are compatible with well-known results from tropical algebraic geometry, such as the Sturmfels–Tevelev multiplicity formula, and indicates how the results generalize from analytifications of algebraic varieties to more general analytic spaces.

The theory of differential forms and currents on analytifications of algebraic varieties was developed in parallel, using rather different methods, by *Boucksom, Favre, and Jonsson*, who provide a survey of their work in this volume. Using their foundational work, they are able to investigate a nonarchimedean analog of the *Monge–Ampère equation* on complex varieties. The uniqueness and existence of solutions in the complex setting are famous theorems of Calabi and Yau, respectively. The nonarchimedean analog of the Monge–Ampère equation was first considered by Kontsevich and Tschinkel, and the uniqueness of solutions (analogous to Calabi’s theorem) was established by Yuan and Zhang. The article of Boucksom, Favre, and Jonsson outlines the authors’ proof of existence in a wide range of cases and concludes with a treatment of the special case of toric varieties.

The survey paper by *Kedlaya* is devoted to another topic of much recent research activity, the radii of convergence of solutions for *p-adic differential equations on curves*. A number of classical results, starting with the work of Dwork and Robba in the 1970s, have recently been improved using a fruitful new point of view, introduced by Baldassarri, based on Berkovich spaces. One studies the radius of convergence as a function on the Berkovich analytification and proves that the behavior of this function is governed by its retraction to a suitable skeleton. Kedlaya discusses the state of the art in this active field, including the recent joint papers of

Poineau and Pulita and the forthcoming work of Baldassari and Kedlaya. He also discusses applications to ramification theory, Artin–Schreier theory, and the Oort conjecture.

The survey paper by *Ducros* gives an introduction to the fundamental recent work of Hrushovski and Loeser on *tameness properties* of the topological spaces underlying Berkovich analytifications. Using model theory, and in particular the theory of *stably dominated types*, Hrushovski and Loeser prove that Berkovich analytifications of algebraic varieties and semi-algebraic sets are locally contractible and have the homotopy type of finite simplicial complexes. (Related results, but with different hypotheses, were proven earlier by Berkovich using completely different methods.) Ducros provides the reader with a gentle introduction to the model theory needed to understand the work of Hrushovski and Loeser.

The research article by *Cartwright* pertains to the general question “What are the possible homotopy types of a Berkovich analytic space?” One way of determining the homotopy type of a Berkovich space is to find a deformation retract onto a *skeleton*, such as the dual complex of the special fiber in a regular semi-stable model. Cartwright has developed a theory of *tropical complexes*, decorating these dual complexes with additional numerical data that makes them behave locally like tropicalizations (so that one can make sense, e.g., of chip-firing moves on divisors in higher dimensions). It is well known that any finite graph can be realized as the dual complex of the special fiber in a regular semi-stable degeneration of curves. Cartwright’s article uses his theory of tropical complexes to prove that a wide range of two-dimensional simplicial complexes, including triangulations of orientable surfaces of genus at least 2, cannot be realized as dual complexes of special fibers of regular semi-stable degenerations.

Tropicalizations of embeddings of algebraic varieties in toric varieties depend on the choice of an embedding. Unless an embedding is chosen carefully, the homotopy type of the analytification might be quite different from that of a given tropicalization. For this reason, one often hunts for *faithful* tropicalizations, in which a fixed skeleton maps homeomorphically onto its image in a manner which preserves the integer affine structure. The article by *Werner* in this volume surveys the state of the art in the hunt for faithful tropicalizations, including Werner’s work with Gubler and Rabinoff generalizing the earlier work of Baker, Payne, and Rabinoff, as well as her work with Hübich and Cueto showing that the tropicalization of the Plücker embedding of the Grassmannian $G(2, n)$ is faithful.

Curves of genus at least 1 over $\mathbb{C}((t))$ have *canonical* minimal skeletons, obtained by taking a minimal regular model over the valuation ring and taking the dual complex of the special fiber. For higher-dimensional varieties, there is no longer a unique minimal regular model. Nevertheless, canonical skeletons do exist in many cases, including for varieties of log-general type (varieties having “sufficiently many pluricanonical forms”). The survey paper by *Nicaise* presents two elegant constructions of this *essential skeleton*, based respectively on Nicaise’s joint work with Mustață and Xu. This work relies crucially on deep facts from the minimal model program and suggests the existence of further relations between birational geometry and the topology of Berkovich spaces yet to be discovered.

The essential skeleton of the analytification of a variety X/K , where K is a discretely valued field, is defined using a certain *weight function* attached to pluricanonical forms. The definition of the weight function uses arithmetic intersection theory and only makes sense over a discretely valued field. The research article by *Temkin* gives a new construction of the essential skeleton which makes sense when K is an *arbitrary* nonarchimedean field and which agrees with the Mustață–Nicaise construction when K is discretely valued of residue characteristic zero. The new construction of Temkin is based on the so-called *Kähler seminorm* on sheaves of relative differential pluriforms. Temkin carefully lays the foundations for the theory of seminorms on sheaves of rings or modules and, as an application, proves generalizations of the main theorems of Mustață and Nicaise.

Both Berkovich’s theory and tropical geometry work equally well over trivially valued fields, but in these cases, one does not have an interesting theory of degenerations to produce skeletons from dual complexes of special fibers. The article by *Abramovich, Chen, Marcus, Ulirsch, and Wise* explains how logarithmic structures on varieties over valued fields produce skeletons of Berkovich analytifications and, moreover, how these skeletons can be endowed with the structure of an *Artin fan*. The authors explain how, following Ulirsch, an Artin fan can be thought of as the nonarchimedean analytification of an Artin stack that locally looks like the quotient of a toric variety by its dense torus. The final section presents a series of intriguing questions for future research.

As mentioned above, in many cases, Berkovich spaces can be understood as limits of tropicalizations. The article by *Foster* gives an expository treatment of recent progress in this direction, presenting joint work with Payne in which the adic analytifications of Huber are realized as limits of *adic tropicalizations*. The underlying topological space of an adic tropicalization is the disjoint union of all initial degenerations. Just as Berkovich spaces are maximal Hausdorff quotients of Huber adic spaces, ordinary tropicalizations are maximal Hausdorff quotients of adic tropicalizations. One technical advantage of adic tropicalizations is that they are *locally ringed spaces* (ordinary tropicalizations do carry a natural structure sheaf, the push-forward of the structure sheaf on the Berkovich analytic space, but the stalks of this sheaf are not local rings).

The wide-ranging survey article of *Baker and Jensen* covers the tropical approach to degenerations of linear series, along with applications to Brill and Noether theory and other problems in algebraic and arithmetic geometry. Starting from Jacobians of graphs, component groups of Néron models, the combinatorics of chip-firing, and tropical geometry of Riemann–Roch, the paper makes connections to Berkovich spaces and their skeletons and also with the classical theory of limit linear series due to Eisenbud and Harris. The concluding sections give overviews of several applications, including the tropical proofs of the Brill–Noether theorem, Gieseker–Petri theorem, and maximal rank conjecture for quadrics, as well as the recent work of Katz, Rabinoff, Zureick, and Brown on uniform bounds for the number of rational points on curves of small Mordell–Weil rank.

The volume ends with the encyclopedic survey article by *Katz*, which provides an introduction to matroid theory aimed at an audience of algebraic geometers.

Highlights of the survey include equivalent descriptions of matroids in terms of matroid polytopes and cohomology classes on the permutahedral toric variety, as well as a discussion of realization spaces and connections to tropical geometry. The article concludes with an exposition of the Huh–Katz proof of Rota’s *log-concavity conjecture* for characteristic polynomials of matroids in the representable case.¹

Atlanta, GA, USA
New Haven, CT, USA

Matthew Baker
Sam Payne

¹While this book was in press, Adiprasito, Huh, and Katz announced a proof of the full Rota conjecture.

Contents

Forms and Currents on the Analytification of an Algebraic Variety (After Chambert-Loir and Ducros)	1
Walter Gubler	
The Non-Archimedean Monge–Ampère Equation	31
Sébastien Boucksom, Charles Favre, and Mattias Jonsson	
Convergence Polygons for Connections on Nonarchimedean Curves	51
Kiran S. Kedlaya	
About Hrushovski and Loeser’s Work on the Homotopy Type of Berkovich Spaces	99
Antoine Ducros	
Excluded Homeomorphism Types for Dual Complexes of Surfaces	133
Dustin Cartwright	
Analytification and Tropicalization Over Non-archimedean Fields	145
Annette Werner	
Berkovich Skeleta and Birational Geometry	173
Johannes Nicaise	
Metrization of Differential Pluriforms on Berkovich Analytic Spaces	195
Michael Temkin	
Skeletons and Fans of Logarithmic Structures	287
Dan Abramovich, Qile Chen, Steffen Marcus, Martin Ulirsch, and Jonathan Wise	
Introduction to Adic Tropicalization	337
Tyler Foster	
Degeneration of Linear Series from the Tropical Point of View and Applications	365
Matthew Baker and David Jensen	

Matroid Theory for Algebraic Geometers 435
Eric Katz

Index 519

List of Contributors

Dan Abramovich Department of Mathematics, Brown University, Providence, RI, USA

Matthew Baker School of Mathematics, Georgia Institute of Technology, Atlanta, GA, USA

Sébastien Boucksom CNRS-CMLS, École Polytechnique, Palaiseau, France

Dustin Cartwright Department of Mathematics, University of Tennessee, Knoxville, TN, USA

Qile Chen Department of Mathematics, Boston College, Boston, MA, USA

Antoine Ducros IMJ-PRG, Paris, France
Institut Universitaire de France, Paris, France

Tyler Foster Department of Mathematics, University of Michigan, Ann Arbor, MI, USA

Charles Favre CNRS-CMLS, École Polytechnique, Palaiseau, France

Walter Gubler Universität Regensburg, Fakultät für Mathematik, Regensburg, Germany

David Jensen Department of Mathematics, University of Kentucky, Lexington, KY, USA

Mattias Jonsson Department of Mathematics, University of Michigan, Ann Arbor, MI, USA

Eric Katz Department of Combinatorics & Optimization, University of Waterloo, Waterloo, ON, Canada

Kiran S. Kedlaya Department of Mathematics, University of California, San Diego, La Jolla, CA, USA

Steffen Marcus Department of Mathematics and Statistics, The College of New Jersey, Ewing, NJ, USA

Johannes Nicaise Department of Mathematics, KU Leuven, Heverlee, Belgium

Michael Temkin Einstein Institute of Mathematics, The Hebrew University of Jerusalem, Giv'at Ram, Jerusalem, Israel

Martin Ulirsch Hausdorff Center for Mathematics, University of Bonn, Bonn, Germany

Annette Werner Institut für Mathematik, Goethe-Universität Frankfurt, Frankfurt, Germany

Jonathan Wise University of Colorado, Boulder, CO, USA

Forms and Currents on the Analytification of an Algebraic Variety (After Chambert-Loir and Ducros)

Walter Gubler

Abstract Chambert-Loir and Ducros have recently introduced real differential forms and currents on Berkovich spaces. In these notes, we survey this new theory and we will compare it with tropical algebraic geometry.

Keywords Arakelov theory • Non-Archimedean geometry • Tropical geometry

MSC2010: 14G22, 14T05

1 Introduction

Antoine Chambert-Loir and Antoine Ducros have recently written the preprint “Formes différentielles réelles et courants sur les espaces de Berkovich” (see [11]). This opens the door for applying methods from differential geometry also at non-Archimedean places. We may think of possible applications for Arakelov theory or for non-Archimedean dynamics. In the Arakelov theory developed by Gillet and Soulé [16], contributions of the p -adic places are described in terms of algebraic intersection theory on regular models over the valuation ring. The existence of such models usually requires the existence of resolution of singularities which is not known in general. Another disadvantage is that canonical metrics of line bundles on Abelian varieties with bad reduction cannot be described in terms of models. In the case of curves, there is an analytic description of Arakelov theory also at finite places due to Chinburgh–Rumely [12], Thuillier [24] and Zhang [26]. Now the paper of Chambert-Loir and Ducros provides us with an analytic formalism including (p, q) -forms, currents and differential operators d', d'' such that the crucial Poincaré–Lelong equation holds. This makes hope that we get also an analytic description of the p -adic contributions in Arakelov theory. In Amaury Thuillier’s thesis [24], he has given a non-Archimedean potential theory on curves. For the case

W. Gubler (✉)

Universität Regensburg, Fakultät für Mathematik, Universitätsstrasse 31,
93040 Regensburg, Germany

e-mail: walter.gubler@mathematik.uni-regensburg.de

of the projective line, we refer to the book of Baker and Rumely [2] with various applications to non-Archimedean dynamics. Again, we may hope to use the paper of Chambert-Loir and Ducros to give generalizations to higher dimensions.

The purpose of the present paper is to summarize the preprint [11] and to compare it with tropical algebraic geometry. We will assume that K is an algebraically closed field endowed with a (nontrivial) complete non-Archimedean absolute value $|\cdot|$. Let $v := -\log|\cdot|$ be the corresponding valuation and let $\Gamma := v(K^\times)$ be the value group. Note that the residue field \bar{K} is also algebraically closed. For the sake of simplicity, we will restrict mostly to the case of an algebraic variety X over K . In this case, there is quite an easy description of the associated analytic space X^{an} and so we require less knowledge about the theory of Berkovich analytic spaces than in [11]. The main idea is quite simple: Suppose that X is an n -dimensional closed subvariety of the split multiplicative torus $T = \mathbb{G}_m^r$. Then there is a tropicalization map $\text{trop} : T^{\text{an}} \rightarrow \mathbb{R}^r$. Roughly speaking, the map is given by applying the valuation v to the coordinates of the points. Tropical geometry says that the tropical variety $\text{Trop}(X) := \text{trop}(X^{\text{an}})$ is a weighted polyhedral complex of pure dimension n satisfying a certain balancing condition. The thesis of Lagerberg [20] gives a formalism of (p, q) -superforms on \mathbb{R}^r together with differential operators d', d'' similar to $\partial, \bar{\partial}$ in complex analytic geometry. Using the tropicalization map, we have a pull-back of these forms and differential operators to X^{an} . In general, we may cover an arbitrary algebraic variety X of pure dimension n by very affine open charts U which means that U has a closed immersion to \mathbb{G}_m^r and we may apply the above to define (p, q) -forms and currents on X^{an} . Chambert-Loir and Ducros prove that there is an integration of compactly supported (n, n) -forms on X^{an} with the formula of Stokes and the Poincaré–Lelong formula. The main result of the paper [11] is that the non-Archimedean Monge–Ampère measures, which were introduced by Chambert-Loir [10] directly as Radon measures on X^{an} , may be written as an n -fold wedge product of first Chern currents. We will focus in this paper on the basics and so we will omit a description of this important result here.

Terminology

In $A \subset B$, A may be equal to B . The complement of A in B is denoted by $B \setminus A$ as we reserve $-$ for algebraic purposes. The zero is included in \mathbb{N} and in \mathbb{R}_+ .

All occurring rings and algebras are with 1. If A is such a ring, then the group of multiplicative units is denoted by A^\times . A variety over a field is an irreducible separated reduced scheme of finite type. We denote by \bar{F} an algebraic closure of the field F .

The terminology from convex geometry is introduced in Sects. 2 and 3. Note that polytopes and polyhedra are assumed to be convex.

This paper was written as a backup for my survey talk at the Simons Symposium on “Non-Archimedean and Tropical Geometry” in St. John from 1.4.2013–5.4.2013. Many thanks to the organizers Matt Baker and Sam Payne for the invitation and the Simons Foundation for the support. The author thanks Antoine Chambert-Loir, Julius Hertel, Klaus Künnemann, Hartwig Mayer, Jascha Smacka and Alejandro Soto for helpful comments. I am very grateful to the referee for his careful reading and his suggestions.

2 Superforms and Supercurrents on \mathbb{R}^r

In this section, we recall the construction of superforms and supercurrents introduced by Lagerberg (see [20, Sect.2]). They are real analogues of complex (p, q) -forms or currents on \mathbb{C}^r . So let us first recall briefly the definitions in complex analytic geometry. On \mathbb{C}^r , we have the holomorphic coordinates z_1, \dots, z_r . A (p, q) -form α is given by

$$\alpha = \sum_{I, J} \alpha_{IJ} dz_I \wedge d\bar{z}_J,$$

where I (resp., J) ranges over all subsets of $\{1, \dots, r\}$ of cardinality p (resp., q) and where the α_{IJ} are smooth functions. Here, use the convenient notation $dz_I := dz_{i_1} \wedge \dots \wedge dz_{i_p}$ and $d\bar{z}_J := d\bar{z}_{j_1} \wedge \dots \wedge d\bar{z}_{j_q}$ for the elements $i_1 < \dots < i_p$ of I and $j_1 < \dots < j_q$ of J . We have linear differential operators d', d'' and $d = d' + d''$ on differential forms which are determined by the rules

$$d'f = \sum_{i=1}^r \frac{\partial f}{\partial z_i} dz_i, \quad d''f = \sum_{j=1}^r \frac{\partial f}{\partial \bar{z}_j} d\bar{z}_j$$

for smooth complex functions f on \mathbb{C}^r . Very often, these differential operators are denoted by $\partial := d'$ and $\bar{\partial} := d''$. A current is a continuous linear functional on the space of differential forms on \mathbb{C}^r . Continuity is with respect to uniform convergence of finitely many derivatives on compact subsets. Differential forms may be viewed as currents using integration and the differential operators d, d', d'' extend to currents. For details, we refer to [14, Chap. I] or to [17].

The goal of this section is to give a real analogue in the following setting: Let N be a free abelian group of rank r with dual abelian group $M := \text{Hom}(N, \mathbb{Z})$. For convenience, we choose a basis e_1, \dots, e_r of N leading to coordinates x_1, \dots, x_r on $N_{\mathbb{R}}$. Our constructions will depend only on the underlying real affine structure and the integration at the end will depend on the underlying integral \mathbb{R} -affine structure, but not on the choice of the coordinates. Here, an *integral \mathbb{R} -affine space* is a real affine space whose underlying real vector space has an integral structure, i.e. it comes with a complete lattice. The definition of the integrals in [11] does use calibrations which makes the integrals in some sense unnatural. In the case of an underlying canonical integral structure (which is the case for tropicalizations), there is a canonical calibration (as in [11, Sect. 3.5]) and both definitions of the integrals are the same.

2.1. Let $A^k(U, \mathbb{R})$ be the space of smooth real differential forms of degree k on an open subset U of $N_{\mathbb{R}}$, then a *superform of bidegree (p, q)* on U is an element of

$$A^{p,q}(U) := A^p(U, \mathbb{R}) \otimes_{C^\infty(U)} A^q(U, \mathbb{R}) = C^\infty(U) \otimes_{\mathbb{Z}} \Lambda^p M \otimes_{\mathbb{Z}} \Lambda^q M.$$

Formally, such a superform α may be written as

$$\alpha = \sum_{|I|=p, |J|=q} \overline{\alpha_{IJ}} d'x_I \wedge d''x_J$$

where I (resp., J) consists of $i_1 < \dots < i_p$ (resp., $j_1 < \dots < j_q$), $\alpha_{IJ} \in C^\infty(U)$ and

$$d'x_I \wedge d''x_J := (dx_{i_1} \wedge \dots \wedge dx_{i_p}) \otimes (dx_{j_1} \wedge \dots \wedge dx_{j_q}).$$

The wedge product is defined in the usual way on the space of superforms $A(U) := \bigoplus_{p,q \leq n} A^{p,q}(U)$ which means that $d'x_i$ and $d''x_j$ anticommute. There is a canonical $C^\infty(U)$ -linear isomorphism $J^{p,q} : A^{p,q}(U) \rightarrow A^{q,p}(U)$ obtained by switching factors in the tensor product. The inverse of $J^{p,q}$ is $J^{q,p}$. We call $\alpha \in A^{p,p}(U)$ *symmetric* if $J^{p,p}\alpha = (-1)^p\alpha$.

2.2. There is a differential operator $d' : A^{p,q}(U) \rightarrow A^{p+1,q}(U)$ given by

$$d'\alpha := \sum_{|I|=p, |J|=q} \sum_{i=1}^r \frac{\partial \overline{\alpha_{IJ}}}{\partial x_i} d'x_i \wedge d'x_I \wedge d''x_J.$$

This does not depend on the choice of coordinates as $d' = d \otimes \text{id}$ on $A^{p,q}(U) = A^p(U, \mathbb{R}) \otimes_{\mathbb{Z}} \wedge^q M$ is an intrinsic characterization using the classical differential d on the space $A^p(U, \mathbb{R})$ of real smooth p -forms. Similarly, we define a differential operator $d'' : A^{p,q}(U) \rightarrow A^{p,q+1}(U)$ by

$$d''\alpha := \sum_{|I|=p, |J|=q} \sum_{j=1}^r \frac{\partial \overline{\alpha_{IJ}}}{\partial x_j} d''x_j \wedge d'x_I \wedge d''x_J.$$

By linearity, we extend these differential operators to $A(U)$. Moreover, we set $d := d' + d''$.

2.3. If N' is a free abelian group of rank r' and if $F : N'_{\mathbb{R}} \rightarrow N_{\mathbb{R}}$ is an affine map with $F(V) \subset U$ for an open subset V of $N'_{\mathbb{R}}$, then we have a well-defined pull-back $F^* : A^{p,q}(U) \rightarrow A^{p,q}(V)$ given as usual. The affine pull-back commutes with the differential operators d , d' and d'' . The pull-back is defined more generally for smooth maps, but then it does not necessarily commute with d , d' and d'' .

2.4. Let $A_c(U)$ denote the space of superforms on U with compact support in U . Recall that r is the rank of M . For $\alpha \in A_c(U)$, we define

$$\int_U \alpha := (-1)^{\frac{r(r-1)}{2}} \int_U \alpha_{LL} dx_1 \wedge \dots \wedge dx_r$$

with $L = \{1, \dots, r\}$ and the usual integration of r -forms with respect to the orientation induced by the choice of coordinates on the right-hand side. If F is an affine map as in 2.3 and if $r = r'$, then we have the transformation formula

$$\int_V F^*(\alpha) = |\det(F)| \int_V \alpha \quad (1)$$

(see [20, Eq. (2.3)]). We conclude that the definition of the integral depends only on the underlying integral \mathbb{R} -affine structure of $N_{\mathbb{R}}$.

The sign $(-1)^{\frac{r(r-1)}{2}}$ is explained by the fact that we want $d'x_1 \wedge d''x_1 \wedge \cdots \wedge d'x_r \wedge d''x_r$ to be a positive (r, r) -superform and hence

$$\int_U f d'x_1 \wedge d''x_1 \wedge \cdots \wedge d'x_r \wedge d''x_r = \int_U f dx_1 \wedge \cdots \wedge dx_r$$

for any $f \in C_c^\infty(U)$ (see [11] for more details about positive forms).

2.5. Now let σ be a *polyhedron* of dimension n in $N_{\mathbb{R}}$. By definition, σ is the intersection of finitely many halfspaces $H_i := \{\omega \in N_{\mathbb{R}} \mid \langle u_i, \omega \rangle \leq c_i\}$ with $u_i \in M_{\mathbb{R}}$ and $c_i \in \mathbb{R}$. A *polytope* is a bounded polyhedron. We say that σ is an *integral G -affine polyhedron* for a subgroup G of \mathbb{R} if we may choose all $u_i \in M$ and all $c_i \in G$. In this case, we have a canonical integral \mathbb{R} -affine structure on the affine space \mathbb{A}_σ generated by σ . If \mathbb{L}_σ is the underlying real vector space of \mathbb{A}_σ , then this integral structure is given by the lattice $N_\sigma := \mathbb{L}_\sigma \cap N$. Using 2.3 and the above, we get a well-defined integral $\int_\sigma \alpha$ for any $\alpha \in A_c^{n,n}(U)$, where U is an open neighbourhood of σ .

2.6. In [11], integration is described in terms of a contraction: Similarly as in differential geometry, we may view a superform $\alpha \in A^{p,q}(U)$ as a multilinear map

$$N_{\mathbb{R}}^{p+q} \longrightarrow C^\infty(U), \quad (n_1, \dots, n_{p+q}) \mapsto \alpha(n_1, \dots, n_{p+q})$$

which is alternating in the variables (n_1, \dots, n_p) and also in $(n_{p+1}, \dots, n_{p+q})$. Let $I \subset \{1, \dots, p+q\}$ be a subset of cardinality s with s' elements contained in $\{1, \dots, p\}$ and hence $s'' = s - s'$ elements in $\{p+1, \dots, p+q\}$. Given vectors $v_1, \dots, v_s \in N_{\mathbb{R}}$, the *contraction* $\langle \alpha; v_1, \dots, v_s \rangle_I \in A^{p-s', q-s''}(U)$ is given by inserting v_1, \dots, v_s for the variables $(n_i)_{i \in I}$ of the above multilinear function.

Using the basis e_1, \dots, e_r of N and assuming $\alpha \in A_c^{r,r}(U)$, the contraction $\langle \alpha; e_1, \dots, e_r \rangle_{\{r+1, \dots, 2r\}}$ is a $(r, 0)$ -superform which may be viewed as a classical r -form on U . Then it is immediately clear from the definitions that we have

$$\int_U \alpha = (-1)^{\frac{r(r-1)}{2}} \int_U \langle \alpha; e_1, \dots, e_r \rangle_{\{r+1, \dots, 2r\}}$$

where we use the usual integration of r -forms on the right. Of course, there is no preference to contract with respect to the last r variables. Similarly, we may view $\langle \alpha; e_1, \dots, e_r \rangle_{\{1, \dots, r\}} \in A_c^{0,r}(U)$ as a classical r -form and we have

$$\int_U \alpha = (-1)^{\frac{r(r-1)}{2}} \int_U \langle \alpha; e_1, \dots, e_r \rangle_{\{1, \dots, r\}}.$$

Next, we are looking for an analogue of Stokes' theorem for superforms.

2.7. Let H be an integral \mathbb{R} -affine halfspace in $N_{\mathbb{R}}$. This means that $H = \{\omega \in N_{\mathbb{R}} \mid \langle u, \omega \rangle \leq c\}$ for some $u \in M$ and $c \in \mathbb{R}$. Using a translation, we may assume that $c = 0$ and hence the boundary ∂H is a linear subspace of $N_{\mathbb{R}}$. Let $[\omega_{\partial H, H}]$ be the generator of $N/(N \cap \partial H) \cong \mathbb{Z}$ which points outwards, i.e. there is $u_{\partial H, H} \in M$ such that $u_{\partial H, H}(H) \leq 0$ and $u_{\partial H, H}(\omega_{\partial H, H}) = 1$. We choose a representative $\omega_{\partial H, H} \in N$ and we note also that $u_{\partial H, H}$ is uniquely determined by the above properties.

2.8. Let U be an open subset of $N_{\mathbb{R}}$ and let σ be an r -dimensional integral \mathbb{R} -affine polyhedron contained in U . For any closed face ρ of codimension 1, let $\omega_{\rho, \sigma} := \omega_{\partial H, H}$ using 2.7 for the affine hyperplane ∂H generated by ρ and the corresponding halfspace containing σ . We note that $\omega_{\rho, \sigma} \in N$ is determined up to addition with elements in $N_{\rho} = N \cap \mathbb{L}_{\rho}$, where \mathbb{L}_{ρ} is the linear hyperplane parallel to ρ .

For $\eta \in A_c^{r-1, r}(U)$, we have introduced the contraction $\langle \eta; \omega_{\rho, \sigma} \rangle_{\{2r-1\}}$ as an element of $A_c^{r-1, r-1}(U)$ which is obtained by inserting the vector $\omega_{\rho, \sigma}$ for the $(2r-1)$ -th argument of the corresponding multilinear function (see 2.6). Note that the restriction of this contraction to ρ does not depend on the choice of the representative $\omega_{\rho, \sigma}$. Then we define

$$\int_{\partial \sigma} \eta := \sum_{\rho} \int_{\rho} \langle \eta; \omega_{\rho, \sigma} \rangle_{\{2r-1\}},$$

where ρ ranges over all closed faces of σ of codimension 1. On the right, we use the integrals of $(r-1, r-1)$ -superforms from 2.4. For $\eta \in A_c^{r, r-1}(U)$, we define similarly

$$\int_{\partial \sigma} \eta := \sum_{\rho} \int_{\rho} \langle \eta; \omega_{\rho, \sigma} \rangle_{\{r\}}.$$

Note that the integrals do depend only on the integral \mathbb{R} -affine structure of $N_{\mathbb{R}}$ but do not depend on the choice of the orientation of $N_{\mathbb{R}}$.

If σ is an integral \mathbb{R} -affine polyhedron of any dimension n and if $\eta \in A_c^{n-1, n}(U)$ for an open subset U of $N_{\mathbb{R}}$ containing σ , then we define $\int_{\partial \sigma} \eta$ by applying the above to the affine space \mathbb{A}_{σ} generated by σ and to the pull-back of η to \mathbb{A}_{σ} . We give now a concrete description of $\int_{\partial \sigma} \eta$ in terms of integrals over classical $(n-1)$ -forms. For every closed face ρ of σ , let $N_{\sigma} = \mathbb{L}_{\sigma} \cap N$ be the canonical integral structure on the affine space generated by σ . If $e_1^{\rho}, \dots, e_{n-1}^{\rho}$ is a basis of N_{ρ} , then $\omega_{\rho, \sigma}, e_1^{\rho}, \dots, e_{n-1}^{\rho}$ is a basis of N_{σ} . We note that the contraction $\langle \eta; \omega_{\rho, \sigma}, e_1^{\rho}, \dots, e_{n-1}^{\rho} \rangle_{\{n, \dots, 2n-1\}}$ may be viewed as a classical $(n-1)$ -form on U and hence we get

$$\int_{\partial \sigma} \eta = \sum_{\rho} \int_{\rho} \langle \eta; \omega_{\rho, \sigma} \rangle_{\{2n-1\}} = (-1)^{\frac{n(n-1)}{2}} \sum_{\rho} \int_{\rho} \langle \eta; \omega_{\rho, \sigma}, e_1^{\rho}, \dots, e_{n-1}^{\rho} \rangle_{\{n, \dots, 2n-1\}}.$$

Proposition 2.9 (Stokes' Formula). *Let σ be an n -dimensional integral \mathbb{R} -affine polyhedron contained in the open subset U of $N_{\mathbb{R}}$. For any $\eta' \in A_c^{n-1, n}(U)$ and any $\eta'' \in A_c^{n, n-1}(U)$, we have*

$$\int_{\sigma} d' \eta' = \int_{\partial \sigma} \eta', \quad \int_{\sigma} d'' \eta'' = \int_{\partial \sigma} \eta''.$$

Proof. This is just a reformulation of Lagerberg [20, Proposition 2.3], in the case of a polyhedron using the formalism introduced above. In the quoted result, the boundary was assumed to be smooth, but as the classical Stokes' formula holds also for polyhedra (see [25, 4.7]), this applies here as well. \square

Proposition 2.10 (Green's Formula). *We consider an n -dimensional integral \mathbb{R} -affine polyhedron σ contained in the open subset U of $N_{\mathbb{R}}$. Assume that $\alpha \in A^{p,p}(U)$ and $\beta \in A^{q,q}(U)$ are symmetric with $p + q = n - 1$ and that the intersection of the supports of α and β is compact. Then we have*

$$\int_{\sigma} \alpha \wedge d' d'' \beta - \beta \wedge d' d'' \alpha = \int_{\partial \sigma} \alpha \wedge d'' \beta - \beta \wedge d'' \alpha.$$

Proof. This follows from Stokes' formula as in [11, Lemma 1.3.8]. \square

2.11. A *supercurrent* on U is a continuous linear functional on $A_c^{p,q}(U)$ where the latter is a locally convex vector space in a similar way as in the classical case. We denote the space of such supercurrents by $D_{p,q}(U)$. As usual, we define the linear differential operators d, d' and d'' on $D(U) := \bigoplus_{p,q} D_{p,q}(U)$ by using $(-1)^{p+q+1}$ times the dual of the corresponding differential operator on $A_c^{p,q}(U)$. The sign is chosen in such a way that the canonical embedding $A^{p,q}(U) \rightarrow D_{r-p,r-q}(U)$ is compatible with the operators d, d' and d'' . Here, $\alpha \in A^{p,q}(U)$ is mapped to $[\alpha] \in D_{r-p,r-q}(U)$ given by $[\alpha](\beta) = \int_{N_{\mathbb{R}}} \alpha \wedge \beta$ for any $\beta \in A_c^{r-p,r-q}(U)$.

3 Superforms on Polyhedral Complexes

We keep the notions from the previous section and we will extend them to the setting of polyhedral complexes. We will introduce tropical cycles and we will characterize them as closed currents of integrations over weighted integral \mathbb{R} -affine polyhedral complexes.

3.1. A *polyhedral complex* \mathcal{C} in $N_{\mathbb{R}}$ is a finite set of polyhedra with the following two properties: Every polyhedron in \mathcal{C} has all its closed faces in \mathcal{C} . If $\Delta, \sigma \in \mathcal{C}$, then $\Delta \cap \sigma$ is a closed face of Δ and σ . Note here that the empty set and also σ are allowed as closed faces of a polyhedron σ (see [19, Appendix A], for details).

A polyhedral complex \mathcal{C} is called *integral G -affine* for a subgroup G of \mathbb{R} if every polyhedron of \mathcal{C} is integral G -affine. The *support* $|\mathcal{C}|$ of \mathcal{C} is the union of all polyhedra in \mathcal{C} . The polyhedral complex \mathcal{C} is called *pure dimensional of dimension n* if every maximal polyhedron in \mathcal{C} has dimension n . We will often use the notation $\mathcal{C}_k := \{\sigma \in \mathcal{C} \mid \dim(\sigma) = k\}$ for $k \in \mathbb{N}$.

3.2. Let \mathcal{C} be a polyhedral complex in $N_{\mathbb{R}}$. A *superform* on \mathcal{C} is the restriction of a superform on (an open subset of) $N_{\mathbb{R}}$ to $|\mathcal{C}|$. This means that two superforms agree if their restrictions to any polyhedron of $|\mathcal{C}|$ agree. Let $A(\mathcal{C})$ be the space of superforms on \mathcal{C} . It is an alternating algebra with respect to the induced wedge product. We have also differential operators d, d' and d'' on $A(\mathcal{C})$ given by restriction of the corresponding operators on $A(N_{\mathbb{R}})$. Let $A^{p,q}(\mathcal{C})$ be the space of (p, q) -superforms on \mathcal{C} . The *support* of $\alpha \in A(\mathcal{C})$ is the complement of $\{\omega \in |\mathcal{C}| \mid \alpha \text{ vanishes identically in a neighbourhood of } \omega\}$ in $|\mathcal{C}|$. We denote by $A_c^{p,q}(\mathcal{C})$ the subspace of $A^{p,q}(\mathcal{C})$ of superforms of compact support.

Let N' be a free abelian group of rank r' and let $F : N'_{\mathbb{R}} \rightarrow N_{\mathbb{R}}$ be an affine map. Suppose that \mathcal{C}' is a polyhedral complex of $N'_{\mathbb{R}}$ with $F(|\mathcal{C}'|) \subset |\mathcal{C}|$, then the pull-back in 2.3 induces a pull-back $F^* : A^{p,q}(\mathcal{C}) \rightarrow A^{p,q}(\mathcal{C}')$.

3.3. A polyhedral complex \mathcal{D} *subdivides* the polyhedral complex \mathcal{C} if they have the same support and if every polyhedron Δ of \mathcal{D} is contained in a polyhedron of \mathcal{C} . In this case, we say that \mathcal{D} is a *subdivision* of \mathcal{C} . All our constructions here will be compatible with subdivisions. This is not a problem for the definition of superforms on \mathcal{C} as they depend only on the support $|\mathcal{C}|$.

A *weight* on a pure dimensional polyhedral complex \mathcal{C} is a function m which assigns to every maximal polyhedron $\sigma \in \mathcal{C}$ a number $m_{\sigma} \in \mathbb{Z}$. Then we get a canonical weight on every subdivision of \mathcal{C} . For a weighted polyhedral complex (\mathcal{C}, m) , only the polyhedra $\Delta \in \mathcal{C}$ which are contained in a maximal dimensional $\sigma \in \mathcal{C}$ with $m_{\sigma} \neq 0$ are of interest. They form a subcomplex \mathcal{D} of \mathcal{C} and we define the support of (\mathcal{C}, m) as the support of \mathcal{D} . The polyhedra of $\mathcal{C} \setminus \mathcal{D}$ will usually be neglected.

3.4. Let (\mathcal{C}, m) be a weighted integral \mathbb{R} -affine polyhedral complex of pure dimension n . For $\alpha \in A_c^{n,n}(\mathcal{C})$, we set

$$\int_{(\mathcal{C}, m)} \alpha := \sum_{\sigma \in \mathcal{C}_n} m_{\sigma} \int_{\sigma} \alpha,$$

where we use integration from 2.4 on the right. We define integrals over the boundary of \mathcal{C} for a superform β in $A_c^{n-1,n}(\mathcal{C})$ or in $A_c^{n,n-1}(\mathcal{C})$ by

$$\int_{\partial(\mathcal{C}, m)} \beta = \sum_{\sigma \in \mathcal{C}_n} m_{\sigma} \int_{\partial\sigma} \beta,$$

where we use the boundary integrals from 2.8 on the right. Note that the boundary $\partial\mathcal{C}$ may be defined as the subcomplex consisting of the polyhedra of dimension at most $n - 1$, but there is no canonical weight on $\partial\mathcal{C}$. Indeed, the boundary integral $\int_{\partial(\mathcal{C}, m)} \beta$ depends on the relative situation $\partial\mathcal{C} \subset \mathcal{C}$ because of the weight m_{σ} and the contraction with respect to the vectors $\omega_{\rho, \sigma}$ used in the definitions. This is similar to the situation in real analysis where boundary integrals depend on the relative orientation. These classical boundary integrals do depend only on the restriction

of the differential form to the boundary which is clearly wrong for our boundary integrals. However, it is still true that $\int_{\partial(\mathcal{C},m)} \beta = 0$ if the support of β is disjoint from $\partial\mathcal{C}$.

Proposition 3.5 (Stokes' Formula). *Let (\mathcal{C}, m) be a weighted integral \mathbb{R} -affine polyhedral complex of pure dimension n . For any $\eta' \in A_c^{n-1,n}(\mathcal{C})$ and any $\eta'' \in A_c^{n,n-1}(\mathcal{C})$, we have*

$$\int_{(\mathcal{C},m)} d'\eta' = \int_{\partial(\mathcal{C},m)} \eta', \quad \int_{(\mathcal{C},m)} d''\eta'' = \int_{\partial(\mathcal{C},m)} \eta''.$$

Proof. This follows immediately from Stokes' formula for polyhedra given in Proposition 2.9. \square

Example 3.6. If (\mathcal{C}, m) is a weighted integral \mathbb{R} -affine polyhedral complex of pure dimension n , then we get a supercurrent $\delta_{(\mathcal{C},m)} \in D_{n,n}(N_{\mathbb{R}})$ by setting $\delta_{(\mathcal{C},m)}(\eta) = \int_{(\mathcal{C},m)} \eta$ for any $\eta \in A_c^{n,n}(N_{\mathbb{R}})$.

3.7. A weighted integral \mathbb{R} -affine polyhedral complex (\mathcal{C}, m) of pure dimension n is called a *tropical cycle* if its weight m satisfies the following *balancing condition*: For every $(n-1)$ -dimensional $\rho \in \mathcal{C}$, we have

$$\sum_{\sigma \in \mathcal{C}_n, \sigma \supset \rho} m_{\sigma} \omega_{\rho,\sigma} \in N_{\rho}.$$

Here, N_{ρ} is the canonical lattice contained in the affine space generated by ρ and $\omega_{\rho,\sigma} \in N_{\sigma}$ is the lattice vector pointing outwards of σ (see 2.8). Tropical cycles are the basic objects in tropical geometry.

Proposition 3.8. *Let (\mathcal{C}, m) be a weighted integral \mathbb{R} -affine polyhedral complex of pure dimension n on $N_{\mathbb{R}}$. Then the following conditions are equivalent:*

- (a) (\mathcal{C}, m) is a tropical cycle;
- (b) $\delta_{(\mathcal{C},m)}$ is a d' -closed supercurrent on $N_{\mathbb{R}}$ and
- (c) $\delta_{(\mathcal{C},m)}$ is a d'' -closed supercurrent on $N_{\mathbb{R}}$.

Proof. Let $\alpha \in A_c^{n-1,n}(N_{\mathbb{R}})$. By Stokes' formula in Proposition 3.5, we have

$$\delta_{(\mathcal{C},m)}(d'\alpha) = \int_{\partial(\mathcal{C},m)} \alpha = \sum_{\rho} \int_{\rho} \langle \alpha; \sum_{\sigma \supset \rho} m_{\sigma} \omega_{\rho,\sigma} \rangle_{\{2n-1\}},$$

where ρ (resp., σ) ranges over all elements of \mathcal{C} of dimension $n-1$ (resp., n). Suppose now that $\sum_{\sigma \supset \rho} m_{\sigma} \omega_{\rho,\sigma} \in N_{\rho}$ for some $(n-1)$ -dimensional $\rho \in \mathcal{C}$. Recall that we may view α as a multilinear map $N_{\mathbb{R}}^{2n-1} \rightarrow C^{\infty}(N_{\mathbb{R}})$ which is alternating in the first $n-1$ arguments and also alternating in the last n arguments. But an alternating n -linear map on a vector space of dimension $n-1$ is zero and hence the restriction of $\langle \alpha; \sum_{\sigma \supset \rho} m_{\sigma} \omega_{\rho,\sigma} \rangle_{\{2n-1\}}$ to ρ is zero. Then the above display proves (a) \Rightarrow (b).

Conversely, if $\sum_{\sigma \supset \rho} m_\sigma \omega_{\rho, \sigma} \notin N_\rho$ for some $(n-1)$ -dimensional $\rho \in \mathcal{C}$, then there is an $\alpha \in A_c^{n-1, n}(N_\mathbb{R})$ such that the restriction of $\langle \alpha; \sum_{\sigma \supset \rho} m_\sigma \omega_{\rho, \sigma} \rangle_{\{2n-1\}}$ to ρ is nonzero. We may also assume that the support of α is disjoint from all other $(n-1)$ -dimensional polyhedra of \mathcal{C} . Then the above display proves (b) \Rightarrow (a). The equivalence of (a) and (c) is shown similarly. \square

3.9. Now let $F : N'_\mathbb{R} \rightarrow N_\mathbb{R}$ be an affine map whose underlying linear map is integral, i.e. induced by a homomorphism $N' \rightarrow N$. We will define the push-forward of a weighted integral \mathbb{R} -affine polyhedral complex (\mathcal{C}', m) of pure dimension n on $N'_\mathbb{R}$. For details, we refer to [1, Sect. 7]. After a subdivision of \mathcal{C}' , we may assume that

$$F_*(\mathcal{C}') := \{F(\sigma') \mid \sigma' \text{ is a face of } v' \in \mathcal{C}' \text{ with } \dim(F(v')) = n\}$$

is a polyhedral complex in $N_\mathbb{R}$. We define the multiplicity of an n -dimensional $F(\sigma') \in F_*(\mathcal{C}')$ by

$$m_{F(\sigma')} := \sum_{v' \in \mathcal{C}'_n, v' \subset F^{-1}(F(\sigma'))} [M'_{v'} : M_{F(\sigma')}] m_{v'}.$$

Endowed with these multiplicities, we get a weighted integral \mathbb{R} -affine polyhedral complex $F_*(\mathcal{C}', m)$ of $N_\mathbb{R}$. If (\mathcal{C}', m) is a tropical cycle, then $F_*(\mathcal{C}', m)$ is also a tropical cycle. It might happen that $F_*(\mathcal{C}', m)$ is empty, then we get the tropical zero cycle.

Proposition 3.10 (Projection Formula). *Using the assumptions above and $\alpha \in A_c^{n, n}(F_*(\mathcal{C}'))$, we have $\int_{F_*(\mathcal{C}', m)} \alpha = \int_{(\mathcal{C}', m)} F^*(\alpha)$.*

Proof. Let σ' be an n -dimensional polyhedron of \mathcal{C}' . Then $\sigma := F(\sigma')$ is an integral \mathbb{R} -affine polyhedron in $N_\mathbb{R}$. We assume for the moment that σ is also n -dimensional. As above, we consider the lattice $N_\sigma := N \cap \mathbb{L}_\sigma$ in $N_\mathbb{R}$, where \mathbb{L}_σ is the linear space which is a translate of the affine space generated by σ . Let A be the matrix of the homomorphism $F : N_{\sigma'} \rightarrow N_\sigma$ with respect to integral bases. Then we have $|\det(A)| = [N_{\sigma'} : N_\sigma]$ and hence the transformation formula (1) shows

$$\int_{\sigma'} F^* \alpha = [N_{\sigma'} : N_\sigma] \int_\sigma \alpha. \quad (2)$$

If $\dim(\sigma) < n$, then both sides are zero and hence formula (2) is true in any case. Using the weighted sum over all σ' , the claim follows immediately from (2). \square

4 Moment Maps and Tropical Charts

A complex manifold is locally defined using analytic charts $\varphi : U \rightarrow \mathbb{C}^r$. The charts help to transport the analysis from \mathbb{C}^r to M . The idea in the non-Archimedean setting is similar to replacing the above charts by algebraic moment maps $\varphi : U \rightarrow \mathbb{G}'_m$

to multiplicative tori and the corresponding tropicalizations $\varphi_{\text{trop}} : U \rightarrow \mathbb{R}^r$. The restriction of φ^{an} to the preimage of an open analytic subset will be called a tropical chart.

In this section, K is an algebraically closed field endowed with a complete nontrivial non-Archimedean absolute value $|\cdot|$. Note that the residue field \tilde{K} is also algebraically closed. Let $v := -\log|\cdot|$ be the associated valuation and let $\Gamma := v(K^\times)$ be the value group. We will study analytifications, tropicalizations and moment maps of the algebraic variety X over K . This will be used in the next section to define (p, q) -forms on X^{an} .

4.1. We recall first the construction of the analytification of X . Let $U = \text{Spec}(A)$ be an open affine subset of X , then U^{an} is the set of multiplicative seminorms on A extending the given absolute value $|\cdot|$ on K . This set is endowed with the topology generated by the functions $U^{\text{an}} \rightarrow \mathbb{R}, p \mapsto p(a)$ with a ranging over A . By glueing, we get a topological space X^{an} which is connected locally compact and Hausdorff. We can endow it with a sheaf of analytic functions leading to a Berkovich analytic space over K which we call the *analytification* of X . For a morphism $\varphi : Y \rightarrow X$ of algebraic varieties over K , we get an analytic morphism $\varphi^{\text{an}} : Y^{\text{an}} \rightarrow X^{\text{an}}$ induced by composing the multiplicative seminorms with φ^\sharp on suitable affine open subsets. We refer to [4] for details, or to [9, Sect. 1.2], for a neat description of the analytification.

4.2. We will define some local invariants in $x \in X^{\text{an}}$. On an open affine neighbourhood $U = \text{Spec}(A)$, the point x is given by a multiplicative seminorm p on A and we often write $|f(x)| := p(f)$ for $f \in A$. Dividing out the prime ideal $I := \{f \in A \mid p(f) = 0\}$, we get a multiplicative norm on the integral domain $B := A/I$ which extends to an absolute value $|\cdot|_x$ on the quotient field of B . The completion of this field is denoted by $\mathcal{H}(x)$. It does not depend on the choice of U and it may be also constructed analytically. The absolute value of $\mathcal{H}(x)$ is denoted by $|\cdot|$ as it extends the given absolute value on K . Note that the completed residue field $\mathcal{H}(x)$ of x remains the same if we replace the ambient variety X by the Zariski closure of x in X .

Let $s(x)$ be the transcendence degree of the residue field of $\mathcal{H}(x)$ over \tilde{K} . The quotient of the value group of $\mathcal{H}(x)$ by Γ is a finitely generated abelian group and we denote its \mathbb{Q} -rank by $t(x)$. Finally, we set $d(x) := s(x) + t(x)$. Note that Abhyankar's inequality shows that $d(x)$ is bounded by the transcendence degree of $\mathcal{H}(x)/K$. By Berkovich [4, Proposition 9.1.3], we have $\dim(X) = \dim(V) = \sup_{x \in V} d(x)$ for every open subset V of X^{an} .

Example 4.3. Let $T = \mathbb{G}_m^r$ be the split multiplicative torus of rank r with coordinates z_1, \dots, z_r . Then a point x of T^{an} could be visualized by the coordinates $z_1(x), \dots, z_r(x) \in \mathcal{H}(x)$ and the multiplicative seminorm corresponding to x is given by $|f(x)| = |f(z_1(x), \dots, z_r(x))|$ for every Laurent polynomial f on T . Conversely, every field extension L/K with an absolute value extending the given absolute value on K and every $(\beta_1, \dots, \beta_r) \in (L^\times)^r$ give rise to a point $x \in T^{\text{an}}$ by $|f(x)| := |f(\beta_1, \dots, \beta_r)|$. Note that L and $(\beta_1, \dots, \beta_r)$ are not uniquely determined by x .

In particular, we get an inclusion of $T(K)$ into T^{an} . For every $x \in T(K)$, we have $d(x) = 0$. However, there can be also other points with $d(x) = 0$. If $T = \mathbb{G}_m^1$, then precisely the points of type 1 (i.e. the K -rational points) and the points of type 4 satisfy $d(x) = 0$ (see [4, 1.4.4]).

Returning to the case $T = \mathbb{G}_m^r$, there are some distinguished points of T^{an} which behave completely different than K -rational points. For positive real numbers s_1, \dots, s_r , we define the *associated weighted Gauss norm* on $K[T]$ by

$$|f|_s := \max_{\mathbf{m} \in \mathbb{Z}^r} |\alpha_{\mathbf{m}}| \mathbf{s}^{\mathbf{m}}$$

for every Laurent polynomial $f = \sum_{\mathbf{m} \in \mathbb{Z}^r} \alpha_{\mathbf{m}} \mathbf{z}^{\mathbf{m}} \in K[T] = K[z_1^{\pm 1}, \dots, z_r^{\pm 1}]$. It follows from the Gauss Lemma that the weighted Gauss norm is a multiplicative seminorm giving rise to a point $\eta_s \in T^{\text{an}}$. The set $S(T^{\text{an}}) := \{\eta_s \mid s_1 > 0, \dots, s_r > 0\}$ is called the *skeleton* of T^{an} . Every point $\eta_s \in S(T^{\text{an}})$ satisfies $d(\eta_s) = r$ (see [15, (0.12) and (0.13)]).

4.4. Let $T := \mathbb{G}_m^r$ be a split multiplicative torus over K with coordinates z_1, \dots, z_r . Then we have the *tropicalization map*

$$\text{trop} : T^{\text{an}} \rightarrow \mathbb{R}^r, \quad p \mapsto (-\log p(z_1), \dots, -\log p(z_r)).$$

It is immediate from the definitions that the map trop is continuous and proper. To get a coordinate free approach, we could use the character group M and its dual N . Then trop is a map from T^{an} to $N_{\mathbb{R}}$. We refer to [19] for details about tropical geometry.

Remark 4.5. Note that we have a natural section $\mathbb{R}^r \rightarrow T^{\text{an}}$ of the tropicalization map. It is given by mapping the point $\omega \in \mathbb{R}^r$ to the weighted Gauss norm η_s associated with $\mathbf{s} := (e^{-\omega_1}, \dots, e^{-\omega_r})$. It follows from [4, Example 5.2.12], that this section is a homeomorphism of \mathbb{R}^r onto a closed subset of T^{an} which is the skeleton $S(T^{\text{an}})$ introduced in 4.3. In this way, we may view the tropicalization map as a map from T^{an} onto $S(X^{\text{an}})$. Then it is shown in [4, Sect. 6.3], that the tropicalization map is a strong deformation retraction of T^{an} onto the skeleton $S(T^{\text{an}})$. This point of view is used very rarely in our paper.

4.6. For a closed subvariety Y of T of dimension n , the *tropical variety associated with Y* is defined by $\text{Trop}(Y) := \text{trop}(Y^{\text{an}})$. The Bieri–Groves theorem says that $\text{Trop}(Y)$ is a finite union of n -dimensional integral Γ -affine polyhedra in \mathbb{R}^r . It is shown in tropical geometry that $\text{Trop}(Y)$ is an integral Γ -affine polyhedral complex. The polyhedral structure is only determined up to subdivision which does not matter for our constructions. We will see below that the tropical variety is endowed with a positive canonical weight m satisfying the balancing condition from 3.7. We get a tropical cycle of pure dimension n which we also denote by $\text{Trop}(Y)$ forgetting the weight m in the notation.

4.7. The *tropical weight* m on an n -dimensional polyhedron σ of $\text{Trop}(Y)$ is defined in the following way. By density of the value group Γ in \mathbb{R} , there is $\omega \in \Gamma^r \cap \text{relint}(\sigma)$. We choose $t \in \mathbb{G}_m^r(K)$ with $\text{trop}(t) = \omega$. Then the closure of $t^{-1}Y$ in $(\mathbb{G}_m^r)_{K^\circ}$ is a flat variety over K° whose special fibre is called the *initial degeneration* $\text{in}_\omega(Y)$ of Y at ω . Note that $\text{in}_\omega(Y)$ is a closed subscheme of $(\mathbb{G}_m^r)_{\tilde{K}}$. Let m_W be the multiplicity of an irreducible component W of $\text{in}_\omega(Y)$. Then the *tropical weight* m_σ is defined by $m_\sigma := \sum_W m_W$, where W ranges over all irreducible components of $\text{in}_\omega(Y)$. One can show that the definition is independent of the choices of ω and t . It is a nontrivial fact from tropical geometry that $(\text{Trop}(Y), m)$ is a tropical cycle (see [19, Sect. 13], for details).

4.8. For an open subset U of the algebraic variety X , a *moment map* is a morphism $\varphi : U \rightarrow T$ to a split multiplicative torus $T := \mathbb{G}_m^r$ over K . The *tropicalization* of φ is

$$\varphi_{\text{trop}} := \text{trop} \circ \varphi^{\text{an}} : U^{\text{an}} \xrightarrow{\varphi^{\text{an}}} T^{\text{an}} \xrightarrow{\text{trop}} \mathbb{R}^r.$$

Obviously, this is a continuous map with respect to the topology on the analytification U^{an} . Note that our moment maps are algebraic which differ from the moment maps in [11] which are defined analytically.

We say that the moment map $\varphi' : U' \rightarrow T'$ of the open subset U' of X *refines* the moment map $\varphi : U \rightarrow T$ if $U' \subset U$ and if there is an affine homomorphism $\psi : T' \rightarrow T$ of the multiplicative tori such that $\varphi = \psi \circ \varphi'$ on U' . Here, an *affine homomorphism* means a group homomorphism composed with a (multiplicative) translation on T . This group homomorphism induces a homomorphism $M \rightarrow M'$ of character lattices. Its dual is the linear part of an integral affine map $\text{Trop}(\psi) : N_{\mathbb{R}} \rightarrow N_{\mathbb{R}}$ such that $\varphi_{\text{trop}} = \text{Trop}(\psi) \circ \varphi'_{\text{trop}}$ on $(U')^{\text{an}}$.

If $\varphi_i : U_i \rightarrow T_i$ are finitely many moment maps of nonempty open subsets U_i of X with $i \in I$, then $U := \bigcap_i U_i$ is a nonempty open subset of X and $\varphi : U \rightarrow \prod_i T_i$, $x \mapsto (\varphi_i(x))_{i \in I}$ is a moment map which refines every φ_i . Moreover, it follows easily from the universal property of the product that every moment map $\varphi' : U' \rightarrow T'$ which refines every φ_i refines also φ .

Lemma 4.9. *Let $\varphi : U \rightarrow \mathbb{G}_m^r$ be a moment map on an open subset U of X and let U' be a nonempty open subset of U . Then $\varphi_{\text{trop}}((U')^{\text{an}}) = \varphi_{\text{trop}}(U^{\text{an}})$.*

Proof. Let $\omega \in \varphi_{\text{trop}}(U^{\text{an}})$. We note that $\varphi_{\text{trop}}^{-1}(\omega)$ is a Laurent domain in U^{an} and hence it has the same dimension as U . We conclude that $\varphi_{\text{trop}}^{-1}(\omega)$ is not contained in the analytification of the lower dimensional Zariski-closed subset $U \setminus U'$ and hence $\omega \in \varphi_{\text{trop}}((U')^{\text{an}})$. \square

4.10. If $f : X_1 \rightarrow X_2$ is a morphism of varieties over K , then we define the *push-forward* of X_1 with respect to f as the cycle $f_*(X_1) := \text{deg}(f)f(X_1)$, where the *degree* of f is defined as $\text{deg}(f) := [K(X_1) : K(f(X_1))]$ if f is generically finite and we set $\text{deg}(f) := 0$ if $[K(X_1) : K(f(X_1))] = \infty$. By restriction, the push-forward can be defined in the same way on prime cycles of X_1 and extends by linearity to all cycles of X_1 .

Now let $\varphi : U \rightarrow T = \mathbb{G}_m^r$ be a moment map of the open subset U of X . By 4.6,

$$\mathrm{Trop}(\varphi_*(U)) := \deg(\varphi)\mathrm{Trop}(\overline{\varphi(U)})$$

is a tropical cycle on \mathbb{R}^r . If φ is generically finite, then this tropical cycle is of pure dimension $\dim(X)$ and the support is equal to $\varphi_{\mathrm{trop}}(U^{\mathrm{an}})$ (see Lemma 4.9).

The following result is called the *Sturmfels–Tevelev multiplicity formula*. It was proved by Sturmfels and Tevelev [22] in the case of a trivial valuation and later generalized by Baker, Payne and Rabinoff [3] for every valued field.

Proposition 4.11. *Let $\varphi' : U' \rightarrow T'$ be a moment map of the nonempty open subset U' of X which refines the moment map $\varphi : U \rightarrow T$, i.e. there is an affine homomorphism $\psi : T' \rightarrow T$ such that $\varphi = \psi \circ \varphi'$ on $U' \subset U$. Then we have*

$$(\mathrm{Trop}(\psi))_*(\mathrm{Trop}(\varphi'_*(U'))) = \mathrm{Trop}(\varphi_*(U))$$

in the sense of tropical cycles (see 3.9).

Proof. In fact, the Sturmfels–Tevelev multiplicity formula is the special case where $X = U'$ is a closed subvariety of T' (see [19, Theorem 13.17], for a proof in our setting deducing it from the original sources). In the general case, we conclude that

$$(\mathrm{Trop}(\psi))_*(\mathrm{Trop}(\varphi'_*(U'))) = \mathrm{Trop}(\psi_*((\varphi')_*(U'))) = \mathrm{Trop}(\varphi_*(U')).$$

Since U' is dense in U , the claim follows. \square

4.12. We will show that every open affine subset U of X has a canonical moment map. We note that the abelian group $M_U := \mathcal{O}(U)^\times / K^\times$ is free of finite rank (see [21, Lemme 1]). Here, we use that K is algebraically closed (or at least that X is geometrically reduced). We choose representatives $\varphi_1, \dots, \varphi_r$ in $\mathcal{O}(U)^\times$ of a basis. This leads to a moment map $\varphi_U : U \rightarrow T_U = \mathrm{Spec}(K[M_U])$. By construction, φ_U refines every other moment map on U . Note that this moment map φ_U is canonical up to (multiplicative) translation by an element of $T_U(K)$.

Let $f : X' \rightarrow X$ be a morphism of algebraic varieties over K and let U' is an open subset of X' with $f(U') \subset U$. Then f^\sharp induces a homomorphism $M_U \rightarrow M_{U'}$ of lattices. We get a canonical affine homomorphism $\psi_{U,U'} : T_{U'} \rightarrow T_U$ of the canonical tori with $\psi_{U,U'} \circ \varphi_{U'} = \varphi_U \circ f$. This will be applied very often in the case where U' is an open subset of U in $X' = X$ and $f = \mathrm{id}$. Then we get a canonical affine homomorphism $\psi_{U,U'} : T_{U'} \rightarrow T_U$.

4.13. Recall that an open subset U of X is called *very affine* if U has a closed embedding into a multiplicative torus. Clearly, the following conditions are equivalent for an open affine subset U of X :

- (a) U is very affine;
- (b) $\mathcal{O}(U)$ is generated as a K -algebra by $\mathcal{O}(U)^\times$;
- (c) the canonical moment map φ_U from 4.12 is a closed embedding.

The intersection of two very affine open subsets is again very affine (see the proof of Proposition 4.16). Moreover, the very affine open subsets of X form a basis for the Zariski topology. We conclude that all local considerations can be done using very affine open subsets.

On a very affine open subset, we will almost always use the canonical moment map $\varphi_U : U \rightarrow T_U$ which is a closed embedding by the above. To simplify the notation, we will set $\text{Trop}(U)$ for the tropical variety of U in T_U . It is a tropical cycle in $(N_U)_{\mathbb{R}}$, where N_U is the dual abelian group of M_U . The tropicalization map will be denoted by $\text{trop}_U := (\varphi_U)_{\text{trop}} : U^{\text{an}} \rightarrow (N_U)_{\mathbb{R}}$. Recall that φ_U is only determined up to translation by an element of $T_U(K)$ and hence trop_U and $\text{Trop}(U)$ are only canonical up to an affine translation. This ambiguity is not a problem as our constructions will be compatible with affine translations.

The following result of Ducros relates the local invariant $d(x)$ from 4.2 with tropical dimensions.

Proposition 4.14. *For $x \in X^{\text{an}}$, there is a very affine open neighbourhood U of x in X such that for any open neighbourhood W of x in the analytic topology of U^{an} , there is a compact neighbourhood V of x in W such that $\text{trop}_U(V)$ is a finite union of $d(x)$ -dimensional integral Γ -affine polytopes.*

Proof. We choose rational functions f_1, \dots, f_s on X with $|f_1(x)| = \dots = |f_s(x)| = 1$ such that the reductions $\tilde{f}_1, \dots, \tilde{f}_s$ form a transcendence basis of the residue field extension of $\mathcal{H}(x)/K$. There are rational functions g_1, \dots, g_t which are regular at x such that $|g_1(x)|, \dots, |g_t(x)|$ form a basis of $(|\mathcal{H}(x)^\times|/|K^\times|) \otimes_{\mathbb{Z}} \mathbb{Q}$. By definition, we have $d(x) = s + t$. By (0.12) in [15], $f_1(x), \dots, f_s(x), g_1(x), \dots, g_t(x)$ reduce to a transcendence basis of the graded residue field extensions of $\mathcal{H}(x)/K$ in the sense of Temkin. There is a very affine open neighbourhood U of x in X such that $f_1, \dots, f_s, g_1, \dots, g_t$ are invertible on U . Let $\varphi_1, \dots, \varphi_r \in \mathcal{O}(U)^\times$ be the coordinates of the canonical moment map $\varphi_U : U \rightarrow T_U = \mathbb{G}_m^r$. Then the graded reductions of $\varphi_1, \dots, \varphi_r$ generate a graded subfield of the graded residue field extension of $\mathcal{H}(x)/K$. By construction, this graded subfield has transcendence degree $d(x)$ over the graded residue field of K . By Ducros [15, Theorem 3.2], $\text{Trop}_U(V)$ is a finite union of integral Γ -affine polytopes for every compact neighbourhood V of x in U^{an} which is strict in the sense of Berkovich [5]. For any open neighbourhood W of x in U^{an} , Theorem 3.4 in [15] shows that there is a compact strict neighbourhood V of x in W such that $\text{trop}_U(V)$ is a finite union of $d(x)$ -dimensional polytopes. \square

4.15. A *tropical chart* (V, φ_U) on X^{an} consists of an open subset V of X^{an} contained in U^{an} for a very affine open subset U of X with $V = \text{trop}_U^{-1}(\Omega)$ for some open subset Ω of $\text{Trop}(U)$. Here the canonical moment map $\varphi_U : U \rightarrow T_U$ from 4.12 plays the role of (tropical) coordinates for V . By 4.13, φ_U is an embedding. The condition $V = \text{trop}_U^{-1}(\Omega)$ means that V behaves well with respect to the tropical coordinates. In particular, $\text{trop}_U(V) = \Omega$ is an open subset of $\text{Trop}(U)$.

We say that the tropical chart $(V', \varphi_{U'})$ is a *tropical subchart* of (V, φ_U) if $V' \subset V$ and $U' \subset U$. We note that the definition of tropical chart here is different from the tropical charts in [11, Sect. 3.1], which consist of an analytic morphism to a split torus and a finite union of polytopes containing the tropicalization.

Proposition 4.16. *The tropical charts on X^{an} have the following properties:*

- (a) *They form a basis on X^{an} , i.e. for every open subset W of X^{an} and for every $x \in W$, there is a tropical chart (V, φ_U) with $x \in V \subset W$. We may find such a V such that the open subset $\text{trop}_U(V)$ of $\text{Trop}(U)$ is relatively compact.*
- (b) *The intersection $(V \cap V', \varphi_{U \cap U'})$ of tropical charts (V, φ_U) and $(V', \varphi_{U'})$ is a tropical subchart of both.*
- (c) *If (V, φ_U) is a tropical chart and if U'' is a very affine open subset of U with $V \subset (U'')^{\text{an}}$, then $(V, \varphi_{U''})$ is a tropical subchart of (V, φ_U) .*

Proof. To prove (a), we may assume that $X = \text{Spec}(A)$ is a very affine scheme. A basis of X^{an} is formed by subsets of the form $V := \{x \in X \mid s_1 < |f_1(x)| < r_1, \dots, s_k < |f_k(x)| < r_k\}$ with all $f_a \in A$ and real numbers $s_a < r_a$. Using the ultrametric triangle inequality as applied to $f_a + \pi$ for a nonzero $\pi \in K$ of small absolute value if $f_a(x) = 0$, it is easy to see that we may choose the basis in such a way that $0 < s_a$ for all $a = 1, \dots, k$. Note that V is contained in the analytification of the very affine open subset $U := \{x \in X \mid f_1(x) \neq 0, \dots, f_k(x) \neq 0\}$ of X . It is obvious that (V, φ_U) is a tropical chart proving (a).

To prove (b), let us consider the moment map

$$\Phi : U \cap U' \rightarrow T_U \times T_{U'}, \quad x \mapsto (\varphi_U(x), \varphi_{U'}(x)).$$

Since X is separated, it is easy to see that Φ is a closed embedding and hence $U \cap U'$ is very affine. By definition of a tropical chart, $\Omega := \text{trop}_U(V)$ (resp., $\Omega' := \text{trop}_{U'}(V')$) is an open subset of $\text{Trop}(U)$ (resp., $\text{Trop}(U')$). Note that

$$\Omega'' := \Phi_{\text{trop}}((U \cap U')^{\text{an}}) \cap (\Omega \times \Omega') \subset (N_U)_{\mathbb{R}} \times (N_{U'})_{\mathbb{R}}$$

is an open subset of $\Phi_{\text{trop}}((U \cap U')^{\text{an}})$. An easy diagram chase yields $\Phi_{\text{trop}}^{-1}(\Omega'') = V \cap V'$. Since $\varphi_{U \cap U'}$ refines the moment map Φ , we deduce that $(V \cap V', \varphi_{U \cap U'})$ is a tropical chart. This proves (b).

Finally, we prove (c). Let $\psi := \psi_{U, U''} : T_{U''} \rightarrow T_U$ be the canonical affine homomorphism from 4.12. Then we have $\text{trop}_U = \text{Trop}(\psi) \circ \text{trop}_{U''}$ on $(U'')^{\text{an}}$. Since (V, φ_U) is a tropical chart, $\Omega := \text{trop}_U(V)$ is an open subset of $\text{Trop}(U)$ and $V = \text{trop}_U^{-1}(\Omega)$. Using $V \subset (U'')^{\text{an}}$, we get $V = \text{trop}_{U''}^{-1}(\Omega'')$ for the open subset $\Omega'' := \text{Trop}(\psi)^{-1}(\Omega)$ of $\text{Trop}(U'')$. We conclude that $(V, \varphi_{U''})$ is a tropical chart proving (c). \square

Remark 4.17. In [11], everything is defined for an arbitrary analytic space. In Sect. 7, we will compare their analytic constructions with our algebraic approach.

5 Differential Forms on Algebraic Varieties

On a complex analytic manifold M , we use open analytic charts $\varphi : U \rightarrow \mathbb{C}^r$ to define (p, q) -forms on U by pull-back. The idea in the non-Archimedean setting is similar to replacing the above charts by tropical charts (V, φ_U) from the previous section in order to pull-back Lagerberg's superforms to U^{an} .

In this section, K is an algebraically closed field endowed with a complete nontrivial non-Archimedean absolute value $|\cdot|$. Let $v := -\log|\cdot|$ be the associated valuation and let $\Gamma := v(K^\times)$ be the value group. The theory could be done for arbitrary fields (see [11]), but it is no serious restriction to assume that K is algebraically closed as the theory is stable under base extension and in the classical setting, the analysis is also done over \mathbb{C} . We will introduce (p, q) -forms on the analytification X^{an} of a n -dimensional algebraic variety X over K .

5.1. We recall from 4.15 that a tropical chart (V, φ_U) consists of an open subset V of U^{an} for a very affine open subset U of X such that $V = \text{trop}_U^{-1}(\Omega)$ for an open subset Ω of $\text{Trop}(U)$. Here, $\varphi_U : U \rightarrow T_U$ is the canonical moment map. It is a closed embedding to the torus $T_U = \text{Spec}(K[M_U])$. The tropical variety $\text{Trop}(U)$ is a tropical cycle of $(N_U)_{\mathbb{R}}$ and $\text{trop}_U : U^{\text{an}} \rightarrow (N_U)_{\mathbb{R}}$ is the tropicalization map. The embedding φ_U is only determined up to translation by an element in $T_U(K)$ and hence the tropical constructions are canonical up to integral Γ -affine isomorphisms.

Suppose that we have another tropical chart $(V', \varphi_{U'})$. Then $(V \cap V', \varphi_{U \cap U'})$ is a tropical chart (see Proposition 4.16) and we get a canonical affine homomorphism $\psi_{U, U \cap U'} : T_{U \cap U'} \rightarrow T_U$ of the underlying tori with $\varphi_U = \psi_{U, U \cap U'} \circ \varphi_{U \cap U'}$ on $U \cap U'$ (see 4.12). The associated affine map $\text{Trop}(\psi_{U, U \cap U'}) : (N_{U \cap U'})_{\mathbb{R}} \rightarrow (N_U)_{\mathbb{R}}$ maps the tropical variety $\text{Trop}(U \cap U')$ onto $\text{Trop}(U)$ (use Lemma 4.9). Then we define the *restriction* of the superform $\alpha \in A^{p,q}(\text{trop}_U(V))$ to a superform $\alpha|_{V \cap V'}$ on $\text{trop}_{U \cap U'}(V \cap V')$ by using the pull-back to $\text{trop}_{U \cap U'}(V \cap V')$ with respect to $\text{Trop}(\psi_{U, U \cap U'})$. This plays a crucial role in the following definition:

Definition 5.2. A *differential form* α of bidegree (p, q) on an open subset V of X^{an} is given by a covering $(V_i)_{i \in I}$ of V by tropical charts (V_i, φ_{U_i}) of X^{an} and superforms $\alpha_i \in A^{p,q}(\text{trop}_{U_i}(V_i))$ such that $\alpha_i|_{V_i \cap V_j} = \alpha_j|_{V_i \cap V_j}$ for every $i, j \in I$. If α' is another differential form of bidegree (p, q) on V given by $\alpha'_j \in A^{p,q}(\text{trop}_{U'_j}(V'_j))$ with respect to the tropical charts $(V'_j, \varphi_{U'_j})_{j \in J}$ covering V , then we consider α and α' as the same differential forms if and only if $\alpha_i|_{V_i \cap V'_j} = \alpha'_j|_{V_i \cap V'_j}$ for every $i \in I$ and $j \in J$. We denote the space of (p, q) -differential forms on V by $A^{p,q}(V)$. As usual, we define the space of differential forms on V by $A(V) := \bigoplus_{p,q} A^{p,q}(V)$. The subspace of differential forms of degree $k \in \mathbb{N}$ is denoted by $A^k(V) := \bigoplus_{p+q=k} A^{p,q}(V)$.

5.3. It is obvious from the definitions that the differential forms form a sheaf on X^{an} . Using the corresponding constructions for superforms on tropical cycles, it is immediate to define the wedge product and differential operators d, d' and d'' on differential forms on V . By 4.6, we have $A^{p,q}(V) = \{0\}$ if $\max(p, q) > \dim(X)$.

For a morphism $\varphi : X' \rightarrow X$ and open subsets V (resp., V') of X^{an} (resp., $(X')^{\text{an}}$) with $\varphi(V') \subset V$, we get a pull-back $\varphi^* : A^{p,q}(V) \rightarrow A^{p,q}(V')$ defined in the following way: Suppose that $\alpha \in A^{p,q}(V)$ is given by the covering $(V_i)_{i \in I}$ and the superforms $\alpha_i \in A^{p,q}(\text{trop}_{U_i}(V_i))$ as above. Then there is a covering $(V'_j)_{j \in J}$ of V' by tropical charts $(V'_j, \varphi_{U'_j})$ which is subordinate to $((\varphi^{\text{an}})^{-1}(V_i))_{i \in I}$. This means that for every $j \in J$, there is $i(j) \in I$ with $V'_j \subset V_{i(j)}$ and $\varphi(U'_j) \subset U_{i(j)}$ for the corresponding very affine open subsets. Then $\varphi^*(\alpha)$ is the differential form on V' given by the covering $(V'_j)_{j \in J}$ and the superforms $\varphi^*(\alpha_{i(j)}) \in A^{p,q}(\text{Trop}(U'_j))$. We leave the details to the reader. This construction is functorial as usual.

Remark 5.4. We obtain the same sheaf of differential forms on X^{an} as in [11, Sect. 3]. In the latter reference, all analytic moment maps were used to define differential forms on X^{an} and so it is clear that our differential forms here are also differential forms in the sense of Chambert-Loir and Ducros [11]. To see the converse, we argue as follows: By Proposition 4.16, tropical charts (V, φ_U) form a basis in X^{an} . It follows from Proposition 7.2 that an analytic moment map $\varphi : V \rightarrow (\mathbb{G}_m^r)^{\text{an}}$ may be locally in $x \in V$ approximated by an algebraic moment map $\varphi' : U' \rightarrow \mathbb{G}_m^r$ such that $(\varphi')_{\text{trop}} = \text{trop} \circ \varphi$ in an open neighbourhood of x in V . Here, U' is a suitable very affine open subset of U with $x \in (U')^{\text{an}}$. It follows from [11, Lemma 3.1.10], that we may use algebraic moment maps to define differential forms in the sense of Chambert-Loir and Ducros [11]. Using that $\varphi_{U'}$ factorizes through φ' (see 4.12), we get the claim.

Definition 5.5. Let α be a differential form on an open subset V of X^{an} . The *support* of α is the complement in V of the set of points x of V which have an open neighbourhood V_x such that $\alpha|_{V_x} = 0$. Let $A_c^{p,q}(V)$ be the space of differential forms of bidegree (p, q) with compact support in V .

Proposition 5.6. Let (V, φ_U) be a tropical chart of X^{an} and let $\alpha \in A^{p,q}(V)$ be given by $\alpha_U \in A^{p,q}(\text{trop}_U(V))$. Then $\alpha = 0$ in $A^{p,q}(V)$ if and only if $\alpha_U = 0$ in $A^{p,q}(\text{trop}_U(V))$.

Proof. See [11, Lemme 3.2.2]. □

Remark 5.7. It follows from Proposition 5.6, that $\text{trop}_U(\text{supp}(\alpha)) = \text{supp}(\alpha_U)$ (see [11, Corollaire 3.2.3]). Note however that not every differential form α on the tropical chart (V, φ_U) is given by a single $\alpha_U \in A^{p,q}(\text{trop}_U(V))$ as in Proposition 5.6.

5.8. In analogy with differential geometry on manifolds, we set $C^{\infty}(V) := A^{0,0}(V)$ for any open subset V of X^{an} and a *smooth function* on V is just a differential form of bidegree $(0, 0)$. Since tropicalization maps are continuous, it is clear that a smooth function is a continuous function on V . By the Stone–Weierstrass theorem, the space $C_c^{\infty}(V)$ of smooth functions with compact support in V is a dense subalgebra of $C_c(V)$ endowed with the supremum norm (see [11, Proposition 3.3.5]).

Definition 5.9. Let $(W_i)_{i \in I}$ be an open covering of an open subset W of X^{an} . A smooth *partition of unity* on W with compact supports subordinated to the covering $(W_i)_{i \in I}$ is a family $(\phi_j)_{j \in J}$ of nonnegative smooth functions with compact support on W with the following properties:

- (i) The family $(\text{supp}(\phi_j))_{j \in J}$ is locally finite on W .
- (ii) We have $\sum_{j \in J} \phi_j \equiv 1$ on W .
- (iii) For every $j \in J$, there is $i(j) \in I$ such that $\text{supp}(\phi_j) \subset W_{i(j)}$.

For the following result, we note that X^{an} is paracompact as it is a σ -compact locally compact Hausdorff space, but not necessarily every open subset of X^{an} is paracompact.

Proposition 5.10. *Let $(W_i)_{i \in I}$ be an open covering of a paracompact open subset W of X^{an} . Then there is a smooth partition of unity $(\phi_j)_{j \in J}$ on W with compact supports subordinated to the covering $(W_i)_{i \in I}$.*

Proof. A locally compact Hausdorff space is paracompact if and only if it is the topological sum of locally compact σ -compact spaces (see [8, Chap. 1, Sect. 9, no. 10, Théorème 5]). Therefore we may assume that W is σ -compact. It is enough to show that for every $x \in W$ and every open neighbourhood V of x in W , there is a nonnegative smooth function ϕ with compact support in V and with $\phi(x) > 0$. Then standard arguments from differential geometry yield the existence of the desired partition of unity (see [25, Theorem 1.11]).

To prove the crucial claim at the beginning of the proof, we may assume that V is coming from a tropical chart (V, φ_U) (see Proposition 4.16). Then $\Omega := \text{trop}_U(V)$ is an open subset of $\text{Trop}(U)$ with $\text{trop}_U^{-1}(\Omega) = V$ and hence there is an open subset $\tilde{\Omega}$ in $(N_U)_{\mathbb{R}}$ with $\Omega = \tilde{\Omega} \cap \text{Trop}(U)$. There is a smooth nonnegative function f on $(N_U)_{\mathbb{R}}$ with compact support in $\tilde{\Omega}$ such that $f(\text{trop}_U(x)) > 0$. Since the tropicalization map is proper, the smooth function $\phi := f \circ \text{trop}_U$ has compact support in V and hence ϕ fulfils the claim. \square

So far, we have seen properties of differential forms which are completely similar to the archimedean case. The next result of Chambert-Loir and Ducros [11, Lemme 3.2.5] shows that the support of a differential form of degree at least one is disjoint from $X(K)$.

Lemma 5.11. *Let W be an open subset of X^{an} . We consider $\alpha \in A^{p,q}(W)$ and $x \in W$ with $d(x) < \max(p, q)$. Then $x \notin \text{supp}(\alpha)$.*

Proof. Using Proposition 4.16 and shrinking the open neighbourhood W of x , we may assume that W is a tropical chart (W, φ_U) on which α is given by the superform $\alpha_U \in A^{p,q}(\text{trop}_U(W))$. By Proposition 4.14, there is a very affine open subset U_x of U and a compact neighbourhood V_x of x in $(U_x)^{\text{an}} \cap W$ such that $\text{trop}_{U_x}(V_x)$ is of dimension $d(x)$. By Proposition 4.16, there is a tropical chart $(V', \varphi_{U'})$ with $x \in V' \subset V_x$ and $U' \subset U_x$. By 4.12, there is an affine homomorphism $\psi : T_{U'} \rightarrow T_U$ such that $\varphi_U = \varphi_{U'} \circ \psi$. Using the same factorization for the tropicalizations, we see that the restriction of α to V' is given by $\text{Trop}(\psi)^*(\alpha_U) \in A^{p,q}(\text{trop}_{U'}(V'))$. The inclusion $U_x \subset U$ yields that trop_{U_x} factorizes through trop_U (use 4.12). Since $V' \subset V_x$, we get $\dim(\text{trop}_U(V')) \leq \dim(\text{trop}_{U_x}(V')) \leq d(x) < \max(p, q)$. As $\text{trop}_U(V') = \text{Trop}(\psi)(\text{trop}_{U'}(V'))$, we conclude that $\text{Trop}(\psi)^*(\alpha_U) = 0$. This proves $\alpha = 0$. \square

Corollary 5.12. *Let W be an open subset of X^{an} and let U be a Zariski open subset of X . If $\alpha \in A^{p,q}(W)$ with $\dim(X \setminus U) < \max(p, q)$, then $\text{supp}(\alpha) \subset W \cap U^{\text{an}}$.*

Proof. Let $x \in W \setminus U^{\text{an}}$. Then 4.2 shows that $d(x) \leq \dim(X \setminus U) < \max(p, q)$. By Lemma 5.11, we get $x \notin \text{supp}(\alpha)$ proving the claim. \square

Proposition 5.13. *Let $\alpha \in A_c^{p,q}(X^{\text{an}})$ be a differential form with $\max(p, q) = \dim(X)$. Then there is a very affine open subset U of X such that $\text{supp}(\alpha) \subset U^{\text{an}}$ and such that α is given on U^{an} by a superform $\alpha_U \in A_c^{p,q}(\text{Trop}(U))$.*

Proof. By assumption, the support of α is a compact subset of X^{an} . We conclude that there are finitely many tropical charts $(V_i, \varphi_{U_i})_{i=1, \dots, s}$ covering $\text{supp}(\alpha)$ such that α is given on V_i by the superform $\alpha_i \in A^{p,q}(\text{trop}_{U_i}(V_i))$. Recall that $\Omega_i := \text{trop}_{U_i}(V_i)$ is an open subset of $\text{Trop}(U_i)$. By 4.13, $U := U_1 \cap \dots \cap U_s$ is a nonempty very affine open subset of X . We define the open subset V of U^{an} by $V := U^{\text{an}} \cap \bigcup_{i=1}^s V_i$. Since $\max(p, q) = \dim(X)$, Corollary 5.12 yields $\text{supp}(\alpha) \subset U^{\text{an}}$. Using 4.12, we see that $\text{trop}_{U_i} = \text{Trop}(\psi_i) \circ \text{trop}_U$ for an affine homomorphism $\psi_i : T_U \rightarrow T_{U_i}$ of tori. Then we have

$$\text{trop}_U(V_i \cap U^{\text{an}}) = (\text{Trop}(\psi_i))^{-1}(\Omega_i) \cap \text{Trop}(U)$$

and we denote this open subset of $\text{Trop}(U)$ by Ω'_i . It follows that the preimage of $\Omega := \bigcup_{i=1}^s \Omega'_i$ with respect to $(\varphi_U)_{\text{trop}}$ is equal to V . We conclude that (V, φ_U) is a tropical chart of X^{an} . Note that α is given on $U^{\text{an}} \cap V_i$ by $\alpha'_i := \text{Trop}(\psi_i)^*(\alpha_i) \in A^{p,q}(\Omega'_i)$. By Proposition 5.6, α'_i agrees with α'_j on $\Omega'_i \cap \Omega'_j$ for every $i, j \in \{1, \dots, s\}$ and hence they define a superform $\alpha_U \in A^{p,q}(\Omega)$. By construction, α_U gives the differential form α on V . It follows from Remark 5.7 that α_U has compact support in Ω . Since α has compact support in V , we conclude that α_U is a superform on $\text{Trop}(U)$ which defines α on U^{an} . \square

5.14. Let $\alpha \in A_c^{n,n}(W)$ for an open subset W of X^{an} , where $n := \dim(X)$. Obviously, we may view α as an (n, n) -form on X^{an} with compact support. We call a very affine open subset U as in Proposition 5.13 a *very affine chart of integration* for α . Then α is given by a superform $\alpha_U \in A_c^{n,n}(\text{Trop}(U))$. We define the *integral of α over W* by

$$\int_W \alpha := \int_{\text{Trop}(U)} \alpha_U.$$

Here, we view $\text{Trop}(U)$ as a tropical cycle (see 4.6) and we integrate as in 3.4.

Lemma 5.15. *For $\alpha \in A_c^{n,n}(W)$, the following properties hold:*

- (a) *If U is a very affine chart of integration for α , then every nonempty very affine open subset U' of U is a very affine chart of integration for α .*
- (b) *The definition of $\int_W \alpha$ is independent of the choice of the very affine chart of integration for α .*

Proof. By Corollary 5.12, $\text{supp}(\alpha) \subset (U')^{\text{an}}$ and (a) follows. To prove (b), it is enough to show

$$\int_{\text{Trop}(U)} \alpha_U = \int_{\text{Trop}(U')} \alpha_{U'} \quad (3)$$

for a nonempty very affine open subset U' of U by using (a). The differential form α is given on U^{an} (resp., $(U')^{\text{an}}$) by $\alpha_U \in A_c^{n,n}(\text{Trop}(U))$ (resp., $\alpha_{U'} \in A_c^{n,n}(\text{Trop}(U'))$). By 4.12, there is an affine homomorphism $\psi : T_{U'} \rightarrow T_U$ of the underlying canonical tori such that $\varphi_U = \text{Trop}(\psi) \circ \varphi_{U'}$. It follows that α is given on U' also by $\text{Trop}(\psi)^*(\alpha_U)$. By Proposition 5.6, we have $\alpha_{U'} = \text{Trop}(\psi)^*(\alpha_U)$. The Sturmfels–Tevelev multiplicity formula shows that $\text{Trop}(\psi)_*(\text{Trop}(U')) = \text{Trop}(U)$ (see Proposition 4.11). Then Proposition 3.10 shows that (3) holds. \square

Proposition 5.16. *Let $\lambda, \rho \in \mathbb{R}$ and let $\alpha, \beta \in A_c^{n,n}(W)$. Then we have*

$$\int_W \lambda\alpha + \rho\beta = \lambda \int_W \alpha + \rho \int_W \beta.$$

Proof. By Lemma 5.15, we may choose a simultaneous very affine chart of integration for both α and β . Then the claim follows by the corresponding property of the integration of superforms. \square

We have also *Stokes' theorem* for differential forms on the open subset W of X^{an} . Note that W has trivial boundary in the algebraic situation [4, Theorem 3.4.1] and hence the boundary does not occur as in the version [11, Theorem 3.12.1] for analytic spaces.

Theorem 5.17. *For $n := \dim(X)$ and $\alpha \in A_c^{2n-1}(W)$, we have $\int_W d'\alpha = \int_W d''\alpha = 0$ and hence $\int_W d\alpha = 0$.*

Proof. By Proposition 5.13, there is a very affine open subset U of X such that $\text{supp}(\alpha) \subset U^{\text{an}}$ and such that α is given on U^{an} by a superform $\alpha_U \in A_c^{2n-1}(\text{Trop}(U))$. Then U is a very affine chart of integration for $d'\alpha$ and $d''\alpha$ and the claim follows from Propositions 3.5 and 3.8. \square

Remark 5.18. Integration of differential forms on complex manifolds is defined by using a partition of unity with compact supports subordinated to a covering by holomorphic charts. Surprisingly, this was not necessary in our non-Archimedean algebraic setting as we have defined integration by using a single suitable tropical chart. In fact, the use of a smooth partition of unity $(\phi_j)_{j \in J}$ with compact supports subordinate to an open covering of W by tropical charts $(V_i, \varphi_{U_i})_{i \in I}$ would not work here directly. To illustrate this, suppose that $\alpha \in A_c^{n,n}(W)$ is given on V_i by $\alpha_i \in A_c^{n,n}(\text{trop}_{U_i}(V_i))$. If the functions ϕ_j are of the form $\phi_j = f_j \circ \text{trop}_{U_{i(j)}}$ for some $V_{i(j)} \supset \text{supp}(\phi_j)$ and $f_j \in C_c^\infty(\text{trop}_{U_{i(j)}}(V_{i(j)}))$, then we could set $\int_W \alpha = \sum_{j \in J} \int_{\text{Trop}(U_{i(j)})} f_j \alpha_{i(j)}$. However, the functions ϕ_j could not be expected to have this form and so this approach fails.

Chambert-Loir and Ducros define integration more generally for differential forms on paracompact good analytic spaces (see [11, Sect. 3.8]). The idea is to use a covering by the interiors of affinoid subdomains. Then there is a smooth partition of unity with supports subordinated to this covering which reduces the problem to defining integration over an affinoid subdomain. But in the affinoid case, one can

find a single tropical chart of integration similarly as in Proposition 5.13. It follows from Remark 7.6 and Proposition 7.11 that both definitions give the same integral on the analytification of an algebraic variety.

6 Currents on Algebraic Varieties

In this section, K is an algebraically closed field endowed with a nontrivial non-Archimedean complete absolute value $|\cdot|$. We consider an open subset W of X^{an} for an algebraic variety X over K of dimension n . Similarly as in the complex case, we will first define a topology on $A_c^{p,q}(W)$ and then we will define currents as continuous linear functionals on this space. We will see that the Poincaré–Lelong equation holds for a rational function.

6.1. Let $(V_i, \varphi_{U_i})_{i \in I}$ be finitely many tropical charts contained in W and let Δ_i be a polytope contained in the open subset $\Omega_i := \text{trop}_{U_i}(V_i)$ of $\text{Trop}(U_i)$. We consider the space $A^{p,q}(V_i, U_i, \Delta_i : i \in I)$ of (p, q) -forms α on W with support in $C := \bigcup_{i \in I} \text{trop}_{U_i}^{-1}(\Delta_i)$ such that α is given on V_i by a superform $\alpha_i \in A^{p,q}(\Omega_i)$ for every $i \in I$. Since the tropicalization map is proper (see 4.4), the set C is compact. Similarly as in the complex case, we endow $A^{p,q}(V_i, U_i, \Delta_i : i \in I)$ with the structure of a locally convex space such that a sequence α_k converges to α if and only if all derivatives of the superforms $\alpha_{k,i}$ converge uniformly to the derivatives of the superform α_i on Δ_i . Here, $\alpha_{k,i}$ (resp., α_i) is the superform on Ω_i which defines α_k (resp., α) on V_i and we mean more precisely the derivatives of the coefficients of $\alpha_{k,i}|_{\Delta_i}$ (resp., $\alpha_i|_{\Delta_i}$). It follows easily from Proposition 4.16 that $A_c^{p,q}(W)$ is the union of all spaces $A^{p,q}(V_i, U_i, \Delta_i : i \in I)$ with $(V_i, U_i, \Delta_i : i \in I)$ ranging over all possibilities as above.

6.2. A *current* on an open subset W of X^{an} is a linear functional T on $A_c^{p,q}(W)$ such that the restriction of T to all subspaces $A^{p,q}(V_i, U_i, \Delta_i : i \in I)$ is continuous. The space of currents is a $C^\infty(W)$ -module denoted by $D_{p,q}(W)$. As usual (cf. 2.11), we define the differential operators d' , d'' and $d := d' + d''$ on the total space of currents $D(W) := \bigoplus_{p,q} D_{p,q}(W)$. Using partitions of unity from Proposition 5.10, one can show that the currents form a sheaf on X^{an} (see [11, Lemme 4.2.5]).

Example 6.3. A signed Radon measure μ on an open subset W of X^{an} induces a current $[\mu] \in D_{0,0}(W)$ by setting $[\mu](f) := \int_{X^{\text{an}}} f d\mu$ using ordinary integration theory on X^{an} . Since the topology on $A_c^{0,0}(W) = C_c^\infty(W)$ is finer than the topology induced by the supremum norm, we conclude that $[\mu]$ is indeed a current on W .

Remark 6.4. Let $\varphi : X' \rightarrow X$ be a proper morphism of algebraic varieties over K . Then there is a linear map $\varphi_* : D_{p,q}((X')^{\text{an}}) \rightarrow D_{p,q}(X^{\text{an}})$, where the *push-forward* $\varphi_*(T') \in D_{p,q}(X^{\text{an}})$ of $T' \in D_{p,q}((X')^{\text{an}})$ is characterized by

$$\varphi_*(T')(\alpha) = T'(\varphi^*(\alpha))$$

for every $\alpha \in A_c^{p,q}(X^{\text{an}})$. It follows from continuity of the map $\varphi^* : A_c^{p,q}(X^{\text{an}}) \rightarrow A_c^{p,q}((X')^{\text{an}})$ that $\varphi_*(T)$ is indeed a current on T . To define the push-forward, we need the fact that a proper algebraic morphism induces a proper morphism between the analytifications which implies that the preimage of a compact subset in X^{an} is compact (see [4, Proposition 3.4.7]).

Example 6.5. We have the *current of integration* $\delta_X \in D_{2n}(X^{\text{an}})$ given by $\delta_X(\alpha) = \int_X \alpha$ for $\alpha \in A_c^{2n}(X^{\text{an}})$. More generally, we define the *current of integration along a closed s -dimensional subvariety Y of X* as the push-forward of $\delta_Y \in D_{2s}(Y^{\text{an}})$ to X^{an} . By abuse of notation, we denote this element of $D_{2s}(X^{\text{an}})$ also by δ_Y . By linearity in the components, we define the current of integration along a cycle on X . If W is an open subset of X^{an} , then we get a current $\delta_W \in D_{2n}(W)$ by restricting δ_X .

6.6. Let $T \in D_{p,q}(W)$ and $\omega \in A^{r,s}(W)$ for an open subset W of X^{an} . Then we define $T \wedge \omega \in D_{p-r,q-s}(W)$ by $(T \wedge \omega)(\alpha) = T(\omega \wedge \alpha)$ for $\alpha \in A_c^{p-r,q-s}(W)$. Since the wedge product with a given form is a continuous operation on $A_c(W)$, it is clear that $T \wedge \omega$ is really a current on W .

Example 6.7. For $\omega \in A^{r,s}(W)$, the *current* $[\omega] \in D_{n-r,n-s}(W)$ associated with ω is defined by $[\omega] := \delta_W \wedge \omega$ and we get an injective linear map $a : A^{r,s}(W) \rightarrow D_{n-r,n-s}(W)$ given by $a(\omega) := [\omega]$.

Proposition 6.8. *Let $\omega \in A_c^{2n}(W)$ for an open subset W of X^{an} . Then there is a unique signed Radon measure μ on W such that $\int_W f d\mu = [\omega](f)$ for every $f \in C_c^\infty(W)$ and we have $|\mu|(W) < \infty$.*

Proof. It is easy to prove that $[\omega]$ induces a continuous linear functional on $C_c^\infty(W)$ where this locally convex vector space is endowed with the subspace topology of $C_c(W)$. By 5.8, this subspace is dense and the Riesz representation theorem proves the claim. \square

6.9. Let us again consider an open subset W of X^{an} . A function $f : W \rightarrow \mathbb{R} \cup \{\pm\infty\}$ is called *locally integrable* if f is integrable with respect to the measure μ associated with any $\omega \in A_c^{2n}(W)$. Then we write $\int_W f \omega := \int_W f d\mu$.

For a locally integrable function f on W and $\eta \in A^{p,q}(W)$, we define $[f \cdot \eta] \in D_{p,q}(W)$ by $[f \cdot \eta](\alpha) := \int_W f \eta \wedge \alpha$ for every $\alpha \in A_c^{n-p,n-q}(W)$.

Chambert-Loir and Ducros proved the *Poincaré–Lelong equation* for rational functions:

Proposition 6.10. *Let f be a rational function on X which is not identically zero. Then $\log |f|$ is a locally integrable function on X^{an} and we have $d'd''[\log |f|] = \delta_{\text{div}(f)}$.*

Proof. See [11, Theorem 4.6.5]. \square

7 Generalizations to Analytic Spaces

The final section shows how our notions fit with the paper [11]. While we are restricted to the algebraic case, the paper of Chambert-Loir and Ducros works for arbitrary analytic spaces. We assume that the reader is familiar with the theory of analytic spaces as given in [5] or [23]. For simplicity, we assume again that K is algebraically closed, endowed with a nontrivial non-Archimedean complete absolute value $|\cdot|$ with corresponding valuation $v := -\log|\cdot|$ and that all occurring analytic spaces are strict in the sense of Berkovich [5]. This situation can always be obtained by base change without changing the theory of differential forms and currents. As usual, we use the value group $\Gamma := v(K^\times)$.

7.1. Let Z be a compact analytic space over K . An *analytic moment map* on Z is an analytic morphism $\varphi : Z \rightarrow T^{\text{an}}$ for a split torus $T = \mathbb{G}_m^r$ over K as before. Let M be the character group of T , then we have $T = \text{Spec}(K[M])$. The map $\varphi_{\text{trop}} := \text{trop} \circ \varphi : Z \rightarrow N_{\mathbb{R}}$ is called the *tropicalization map* of φ and we may use the coordinates on T to identify $N_{\mathbb{R}}$ with \mathbb{R}^r .

The next result shows that for the construction of differential forms in the algebraic case, we may restrict our attention to algebraic moment maps.

Proposition 7.2. *Let X be an algebraic variety over K and let $\varphi : W \rightarrow T^{\text{an}}$ be an analytic moment map defined on an open subset W of X^{an} . For every $x \in W$, there is a very affine open subset U of X with an algebraic moment map $\varphi' : U \rightarrow T$ and an open neighbourhood V of x in $U^{\text{an}} \cap W$ such that $\varphi_{\text{trop}} = \varphi'_{\text{trop}}$ on V .*

Proof. We may assume that $X = \text{Spec}(A)$. Similarly as in the proof of Proposition 4.16, there is a neighbourhood $V' := \{x \in X \mid s_1 \leq |f_1(x)| \leq r_1, \dots, s_k \leq |f_k(x)| \leq r_k\}$ of x in W with all $f_a \in A$ and real numbers $0 < s_a < r_a$. We may assume that f_1, \dots, f_k form an affine coordinate system y_1, \dots, y_k on X . Using coordinates on $T = \mathbb{G}_m^r$, the moment map φ is given by analytic functions $\varphi_1, \dots, \varphi_r$ on W which restrict to strictly convergent Laurent series in y_1, \dots, y_k on V' . Cutting the Laurent series in sufficiently high positive and negative degree, we get Laurent polynomials p_1, \dots, p_r with $|p_a| = |\varphi_a|$ on V' for $a = 1, \dots, r$. By Proposition 4.16, there is a very affine open subset U of X such that U^{an} contains x and such that p_1, \dots, p_r define an algebraic moment map $\varphi' : U \rightarrow T$ with $\varphi_{\text{trop}} = \varphi'_{\text{trop}}$ on V' . Choosing a neighbourhood V of x in $U^{\text{an}} \cap V'$, we get the claim. \square

We have the following generalization of the Bieri–Groves theorem. Working with analytic spaces, the boundary ∂Z of Z becomes an issue.

Theorem 7.3 (Berkovich, Ducros). *If Z is a compact analytic space over K of dimension n and if $\varphi : Z \rightarrow T^{\text{an}}$ is an analytic moment map, then $\varphi_{\text{trop}}(Z)$ is a finite union of integral Γ -affine polytopes of dimension at most n . Moreover, $\varphi_{\text{trop}}(\partial Z)$ is contained in a finite union of integral Γ -affine polytopes of dimension $\leq n - 1$. If Z is affinoid, then $\varphi_{\text{trop}}(\partial Z)$ is equal to a finite union of such polytopes.*

Proof. The first claim is due to Berkovich and the remaining claims are due to Ducros (see [15, Theorem 3.2]). \square

7.4. We consider now a compact analytic space Z over K of pure dimension n . Theorem 7.3 shows that the *tropical variety* $\varphi_{\text{trop}}(Z)$ is the support of an integral Γ -affine polytopal complex in $N_{\mathbb{R}}$. Our next goal is to endow this complex with canonical tropical multiplicities. This will lead to the definition of a weighted polytopal complex $(\varphi_{\text{trop}})_*(\text{cyc}(Z))$ which is canonical up to subdivision.

If $\dim(\varphi_{\text{trop}}(Z)) < n$, then we set $(\varphi_{\text{trop}})_*(\text{cyc}(Z)) = 0$ meaning that we choose all tropical weights equal to zero. It remains to consider the case $\dim(\varphi_{\text{trop}}(Z)) = n$. We choose a generic surjective homomorphism $q : T \rightarrow T'$ onto a split multiplicative torus $T' = \text{Spec}(K[M'])$ of rank $n = \dim(Z)$. Generic means that the corresponding linear map $F := \text{Trop}(q)$ is injective on every polytope contained in $\varphi_{\text{trop}}(Z)$. By Theorem 7.3, there is an integral Γ -affine polytopal complex \mathcal{C} in $N_{\mathbb{R}}$ with $|\mathcal{C}| = \varphi_{\text{trop}}(Z)$ such that $F(\tau)$ is disjoint from $(q \circ \varphi)_{\text{trop}}(\partial Z)$ for every n -dimensional face σ of \mathcal{C} and $\tau := \text{relint}(\sigma)$. By passing to a subdivision, we may assume that $F_*(\mathcal{C})$ is a polyhedral complex in $N'_{\mathbb{R}}$ as in 3.9, where N' is the dual of M' as usual.

We identify $F(\tau) \subset N'_{\mathbb{R}} \cong \mathbb{R}^n$ with an open subset of the skeleton $S((T')^{\text{an}})$ as in Remark 4.5. Then it is clear that $q \circ \varphi$ restricts to a map $(q \circ \varphi)^{-1}(\tau) \rightarrow F(\tau)$ which agrees with $F \circ \varphi_{\text{trop}}$ using the identification $S((T')^{\text{an}}) = N'_{\mathbb{R}}$. It is shown in [11, Sect. 2.4], that this restriction of $q \circ \varphi$ is a finite flat and surjective morphism which means that every point p of $F(\tau)$ has a neighbourhood W' in $(T')^{\text{an}}$ such that $(q \circ \varphi)^{-1}(W') \rightarrow W'$ has these properties. Using that $F^{-1}(F(\tau)) = \coprod_{\tau'} \tau'$, where τ' is ranging over all open faces of \mathcal{C} with $F(\tau') = F(\tau)$, we get

$$(q \circ \varphi)^{-1}(F(\tau)) = \coprod_{\tau'} \varphi_{\text{trop}}^{-1}(\tau') \cap (q \circ \varphi)^{-1}(F(\tau)).$$

We conclude that the map $\varphi_{\text{trop}}^{-1}(\tau) \cap (q \circ \varphi)^{-1}(F(\tau)) \rightarrow F(\tau)$ is finite, flat and surjective. Again, this has to be understood in some open neighbourhoods. Since τ is connected, the corresponding degree depends only on τ and not on the choice of p . We denote this degree by $[\varphi_{\text{trop}}^{-1}(\tau) : F(\tau)]$.

Recall that N_{σ} is the canonical lattice in the affine space generated by σ . Then the character lattice M' of T' is of finite index in $M_{\sigma} = \text{Hom}(N_{\sigma}, \mathbb{Z})$.

Definition 7.5. Using the notation from above, the *tropical multiplicity* m_{σ} along σ is defined by

$$m_{\sigma} := [\varphi_{\text{trop}}^{-1}(\tau) : F(\tau)] \cdot [M_{\sigma} : M']^{-1}.$$

Furthermore, $(\varphi_{\text{trop}})_*(\text{cyc}(Z))$ is the weighted polyhedral complex \mathcal{C} endowed with these tropical multiplicities. The weights might be rational numbers, at least we have no argument that they are integers in the analytic case.

Remark 7.6. It is not so easy to show that the tropical multiplicity is well-defined, i.e. independent of the choice of q . Chambert-Loir and Ducros do not use tropical multiplicities, but the latter are equivalent to the canonical calibration introduced in [11, Sect. 3.5]. To summarize this construction, let e_1, \dots, e_n (resp., f_1, \dots, f_n) be a basis of M' (resp., M_σ). Then the canonical calibration of σ is defined as

$$[\varphi_{\text{trop}}^{-1}(\tau) : F(\tau)] \cdot (F|_{(N_\sigma)_\mathbb{R}})^*(e_1 \wedge \dots \wedge e_n) \in \Lambda^n((N_\sigma)_\mathbb{R})$$

together with the orientation induced by the pull-back of e_1, \dots, e_n with respect to the linear isomorphism $F|_{(N_\sigma)_\mathbb{R}}$. The canonical calibration is equal to the calibration $m_\sigma f_1 \wedge \dots \wedge f_n$ together with the orientation induced by f_1, \dots, f_n . Since the canonical calibration does not depend on the choice of q up to refinement [11, Sect. 3.5], the same is true for the tropical multiplicities.

Remark 7.7. One can define the irreducible components of an analytic space (see [13]). A compact analytic space Z has finitely many irreducible components Z_i . Then we define the cycle $\text{cyc}(Z)$ associated with Z as a positive formal \mathbb{Z} -linear combination of the irreducible components Z_i by restriction to affinoid subdomains and then by glueing (see [18, Sect. 2]). One can show that the weighted n -dimensional polyhedral complex $(\varphi_{\text{trop}})_*(\text{cyc}(Z))$ depends only on $\text{cyc}(Z)$ and this dependence is linear. We leave the details to the reader.

The next result shows that the Sturmfels–Tevelev multiplicity formula holds for analytic spaces.

Proposition 7.8. *Let Z be a compact analytic space over K of pure dimension n , let $\varphi : Z \rightarrow T^{\text{an}}$ be an analytic moment map and let $\psi : T \rightarrow T'$ be an affine homomorphism of tori. Then we have*

$$\text{Trop}(\psi)_*((\varphi_{\text{trop}})_*(\text{cyc}(Z))) = ((\psi \circ \varphi)_{\text{trop}})_*(\text{cyc}(Z)).$$

Proof. The corresponding statement for canonical calibrations is shown in [11, Lemma 3.5.2], and hence the claim follows from Remark 7.6. \square

Proposition 7.9. *Let Z be a compact analytic space over K of pure dimension n and let \mathcal{C} be the same integral Γ -affine polytopal complex with support $\varphi_{\text{trop}}(Z)$ as in 7.4. Then for every $(n - 1)$ -dimensional polyhedron ρ of \mathcal{C} not contained in $\varphi_{\text{trop}}(\partial Z)$, the balancing condition*

$$\sum_{\sigma \in \mathcal{C}_n, \sigma \supset \rho} m_\sigma \omega_{\rho, \sigma} \in N_\rho$$

from 3.7 holds in ρ .

Proof. Chambert-Loir and Ducros prove in [11, Theorem 3.6.1], that ρ is harmonious in \mathcal{C} which is a condition for the canonical calibration equivalent to the balancing condition by Remark 7.6. \square

7.10. In an algebraic setting, our goal is to compare the tropical multiplicities introduced in 4.7 with the ones from Definition 7.5. Let us consider an algebraic variety X over K of dimension n and an algebraic moment map $\varphi : X \rightarrow T = \text{Spec}(K[M]) \cong \mathbb{G}_m^r$ over K . Note that $\varphi_{\text{trop}}(X^{\text{an}}) = \text{Trop}(\varphi(X))$. We endow $\varphi_{\text{trop}}(X^{\text{an}})$ with the tropical multiplicities m_σ^{alg} of the tropical cycle $\text{Trop}(\varphi_*(X)) := \text{deg}(\varphi)\text{Trop}(\overline{\varphi(X)})$ of $N_{\mathbb{R}}$.

The analytification X^{an} is not compact (unless $n = 0$), but as $\partial X = \emptyset$, we can define tropical multiplicities in the same analytic manner as in Definition 7.5. This means that we choose a generic projection $q : T \rightarrow T' = \text{Spec}(K[M'])$ onto a torus T' of rank n and an integral Γ -affine polyhedral complex \mathcal{C} with support equal to $\varphi_{\text{trop}}(X^{\text{an}})$ such that $F_*(\mathcal{C})$ is a polyhedral complex on $N'_{\mathbb{R}}$ for the associated linear map $F : N_{\mathbb{R}} \rightarrow N'_{\mathbb{R}}$. For every $\sigma \in \mathcal{C}_n$ and $\tau := \text{relint}(\sigma)$, we define

$$m_\sigma^{\text{an}} := [\varphi_{\text{trop}}^{-1}(\tau) : F(\tau)] \cdot [M_\sigma : M']^{-1}$$

as in Definition 7.5. Since the tropical multiplicities m^{alg} and m^{an} are compatible with subdivision, we may assume that the underlying integral Γ -affine polyhedral complex \mathcal{C} is the same in both definitions.

Now we are ready to compare these two tropical multiplicities.

Proposition 7.11. *Let $\varphi : X \rightarrow T$ be an algebraic moment map and $n := \dim(X)$. Using the notations from above, we have $m_\sigma^{\text{an}} = m_\sigma^{\text{alg}}$ for every $\sigma \in \mathcal{C}_n$.*

Proof. The following argument is quite close to the proof of the Sturmfels–Tevelev formula given by Baker, Payne and Rabinoff (see [3, Theorem 8.2]). We may assume that φ is generically finite, otherwise $\varphi_{\text{trop}}(X^{\text{an}})$ has dimension $< n$ and all tropical multiplicities are zero. Let Y be the closure of $\varphi(X)$ in T and let $q : T \rightarrow T'$ be a generic homomorphism onto a split torus $T' = \text{Spec}(K[M'])$ of rank n with associated linear map $F : N_{\mathbb{R}} \rightarrow N'_{\mathbb{R}}$. There is an open dense subset U of Y such that φ is finite over U . Since the tropical multiplicities m^{an} and m^{alg} are compatible with subdivision of the polyhedral complex \mathcal{C} , we may assume that $\text{Trop}(Y \setminus U)$ is contained in the support of \mathcal{C}_{n-1} .

Let Y' be the closure of $q(Y)$ in T' and let $\omega \in \tau \cap N_\Gamma$. We consider the affinoid subdomains $U_\omega := \text{trop}^{-1}(\omega)$ in T^{an} and $U'_\omega := \text{trop}^{-1}(\omega')$ in $(T')^{\text{an}}$. By finiteness of φ over U , the set $X_\omega := (\varphi^{\text{an}})^{-1}(U_\omega) = \varphi_{\text{trop}}^{-1}(\omega)$ is an affinoid subdomain of X^{an} and φ restricts to a finite morphism $X_\omega \rightarrow Y_\omega := Y^{\text{an}} \cap U_\omega$. Let $\mathcal{X}_\omega, \mathcal{Y}_\omega, \mathcal{U}_\omega, \mathcal{U}'_\omega$ be the canonical formal affine K° -models of $X_\omega, Y_\omega, U_\omega, U'_\omega$ associated with the algebra of power bounded elements in the corresponding affinoid algebra. Moreover, let \overline{Y}_ω be the closure of Y_ω in \mathcal{U}_ω . Then we have canonical morphisms

$$\mathcal{X}_\omega \xrightarrow{\varphi} \mathcal{Y}_\omega \xrightarrow{\iota} \overline{Y}_\omega \xrightarrow{q} \mathcal{U}'_\omega \quad (4)$$

of admissible formal affine schemes over K° in the sense of Bosch, Lütkebohmert and Raynaud (see [6, Sect. 1]). We claim that all these morphisms are finite and

surjective. Obviously, the generic fibres of the first and second morphism are finite and surjective. To see that the generic fibre of the third morphism is finite, we note first that $F^{-1}(\omega') \cap \text{Trop}(Y)$ is finite by construction of q and hence $q^{-1}(U'_{\omega'}) \cap Y^{\text{an}}$ is in the relative interior of an affinoid subdomain of T^{an} which is contained in $q^{-1}(U'_{\omega'})$. We conclude that $q^{-1}(U'_{\omega'}) \cap Y^{\text{an}} \rightarrow U'_{\omega'}$ is a proper map (see the proof of Theorem 4.31 in [3] for more details about the argument). Since $q^{-1}(U'_{\omega'}) \cap Y^{\text{an}}$ is the disjoint union of the finitely many affinoids $U_{\rho} \cap Y^{\text{an}}$, $\rho \in F^{-1}(\omega') \cap \text{Trop}(Y)$, we conclude that q induces a proper morphism $Y_{\omega} \rightarrow U'_{\omega'}$ of affinoids. By Kiehl's direct image theorem [7, Theorem 9.6.3/1], this morphism is finite and hence also surjective using dimensionality arguments. We conclude that all three morphisms in (4) are surjective and finiteness follows from [3, Proposition 3.13].

The degree $[X_{\omega} : U'_{\omega'}]$ of X_{ω} over the affinoid torus $U'_{\omega'}$ is well-defined as $U'_{\omega'}$ is irreducible (see [3, Sect. 3] for a discussion of degrees). Since the degree does not change by passing to an affinoid subdomain of $U'_{\omega'}$ (see [3, Proposition 3.30]), we get

$$[\varphi_{\text{trop}}^{-1}(\tau) : F(\tau)] = [X_{\omega} : U'_{\omega'}]. \quad (5)$$

The projection formula [3, Proposition 3.32] shows

$$[X_{\omega} : U'_{\omega'}] = \sum_B [B : (\mathcal{X}'_{\omega'})_s] = \sum_B [B : (\mathbb{G}_m^n)_{\bar{K}}], \quad (6)$$

where B ranges over all irreducible components of $(\mathcal{X}'_{\omega'})_s$. We conclude from (5) and (6) that

$$[\varphi_{\text{trop}}^{-1}(\tau) : F(\tau)] = \sum_C \sum_{B \text{ over } C} [B : C] \cdot [C : (\mathbb{G}_m^n)_{\bar{K}}], \quad (7)$$

where C ranges over all irreducible components of $(\overline{Y_{\omega}})_s$ and B ranges over all irreducible components of $(\mathcal{X}'_{\omega'})_s$ mapping onto C . Since the special fibre of $\overline{Y_{\omega}}$ is isomorphic to the initial degeneration $\text{in}_{\omega}(Y)$, all irreducible components C are isomorphic to the torus $\text{Spec}(\bar{K}[M_{\sigma}])$ (see [3, Theorem 4.29]) proving

$$[C : (\mathbb{G}_m^n)_{\bar{K}}] = [M_{\sigma} : M']. \quad (8)$$

Using (7) and (8), we get

$$m_{\sigma}^{\text{an}} = [\varphi_{\text{trop}}^{-1}(\tau) : F(\tau)] \cdot [M_{\sigma} : M']^{-1} = \sum_C \sum_{B \text{ over } C} [B : C]. \quad (9)$$

Since X_{ω} is the preimage of the affinoid subdomain Y_{ω} of T^{an} , we deduce from [3, Proposition 3.30], that X_{ω} is of degree $\deg(\varphi)$ over Y_{ω} and hence the projection formula again shows the equality

$$\deg(\varphi) \text{cyc}((\overline{Y_{\omega}})_s) = (\iota \circ \varphi)_*(\text{cyc}((\mathcal{X}'_{\omega'})_s)) \quad (10)$$

of cycles in $(\mathcal{Y}_\omega)_s$. Inserting (10) in (9) by using that the special fibre of \mathcal{X}_ω is reduced, we get

$$m_\sigma^{\text{an}} = \deg(\varphi) \sum_C m(C, (\overline{Y_\omega})_s),$$

where $m(C, (\overline{Y_\omega})_s)$ is the multiplicity of the irreducible component C in the special fibre of $\overline{Y_\omega}$. By definition, the right-hand side is equal to m_σ^{alg} which proves the claim. \square

Remark 7.12. Note that in the algebraic case, Proposition 7.11 yields that the tropical multiplicities in Definition 7.5 are well-defined integers, i.e. independent of the choice of the generic projection q . Moreover, the argument of Chambert-Loir and Ducros for Proposition 7.9 gives a new proof for the classical balancing condition for tropical varieties which is based mainly on degree considerations.

References

1. Allermann, L., Rau, J.: First steps in tropical intersection theory. *Math. Z.* **264**(3), 633–670 (2010)
2. Baker, M., Rumely, R.: *Potential Theory and Dynamics on the Berkovich Projective Line*. Mathematical Surveys and Monographs, vol. 159. American Mathematical Society, Providence, RI (2010)
3. Baker, M., Payne, S., Rabinoff, J.: Nonarchimedean geometry, tropicalization, and metrics on curves. (2011) arXiv:1104.0320v2
4. Berkovich, V.G.: *Spectral Theory and Analytic Geometry over Non-Archimedean Fields*. Mathematical Surveys and Monographs, vol. 33. American Mathematical Society, Providence, RI (1990)
5. Berkovich, V.G.: Étale cohomology for non-Archimedean analytic spaces. *Publ. Math. IHES* **78**, 5–161 (1993)
6. Bosch, S., Lütkebohmert, W.: Formal and rigid geometry. I: rigid spaces. *Math. Ann.* **295**(2), 291–317 (1993)
7. Bosch, S., Güntzer, U., Remmert, R.: *Non-Archimedean Analysis. A Systematic Approach to Rigid Analytic Geometry*. Grundlehren der Mathematischen Wissenschaften, vol. 261. Springer, Berlin (1984)
8. Bourbaki, N.: *Éléments de mathématique: Topologie générale*. Chapitres 1 à 4. Hermann, Paris (1971)
9. Burgos Gil, J.I., Philippon, P., Sombra, M.: Arithmetic geometry of toric varieties. *Metrics, measures and heights*. *Astérisque* **360** (2014)
10. Chambert-Loir, A.: Mesure et équidistribution sur les espaces de Berkovich. *J. Reine Angew. Math.* **595**, 215–235 (2006)
11. Chambert-Loir, A., Ducros, A.: Formes différentielles réelles et courants sur les espaces Berkovich. arXiv:1204.64591
12. Chinburg, T., Rumely, R.: The capacity pairing. *J. Reine Angew. Math.* **434**, 1–44 (1993)
13. Conrad, B.: Irreducible components of rigid spaces. *Ann. Inst. Fourier (Grenoble)* **49**(2), 473–541 (1999)
14. Demailly, J.-P.: *Complex Analytic and Differential Geometry*. Version from 21 June 2012, available at <https://www-fourier.ujf-grenoble.fr/~demailly/manuscripts/agbook.pdf> (2012)

15. Ducros, A.: Espaces de Berkovich, polytopes, squelettes et théorie des modèles. *Confluentes Math.* **4**(4), 1250007, pp. 57 (2012)
16. Gillet, H., Soulé, C.: Arithmetic intersection theory. *Publ. Math. IHES* **72**, 94–174 (1990)
17. Griffiths, P., Harris, J.: *Principles of Algebraic Geometry*. Pure and Applied Mathematics. Wiley, New York (1978)
18. Gubler, W.: Local heights of subvarieties over non-Archimedean fields. *J. Reine Angew. Math.* **498**, 61–113 (1998)
19. Gubler, W.: A guide to tropicalizations. In: *Algebraic and Combinatorial Aspects of Tropical Geometry*. Contemporary Mathematics, pp. 125–189, vol. 589. American Mathematical Society, Providence, RI (2013)
20. Lagerberg, A.: Super currents and tropical geometry. *Math. Z.* **270**, 1011–1050 (2012)
21. Samuel, P.: À propos du théorème des unités. *Bull. Sci. Math.* **90**(2), 89–96 (1966)
22. Sturmfels, B., Tevelev, J.: Elimination theory for tropical varieties. *Math. Res. Lett.* **15**(2–3), 543–562 (2008)
23. Temkin, M.: Introduction to Berkovich analytic spaces. In: *Berkovich Spaces and Applications*. Lecture Notes in Mathematics, vol. 2119, pp. 3–66, Springer, Cham (2015)
24. Thuillier, A.: *Théorie du potentiel sur les courbes en géométrie analytique non archimédienne. Applications à la théorie d’Arakelov*. Thèse de l’Université de Rennes 1 (2005)
25. Warner, F.: *Foundations of Differentiable Manifolds and Lie Groups*. Graduate Texts in Mathematics, vol. 94. Springer, New York (1983)
26. Zhang, S.: Admissible pairing on a curve. *Invent. Math.* **112**, 171–193 (1993)

The Non-Archimedean Monge–Ampère Equation

Sébastien Boucksom, Charles Favre, and Mattias Jonsson

Abstract We give an introduction to our work on the solution to the non-Archimedean Monge–Ampère equation and make comparisons to the complex counterpart. These notes are partially based on talks at the 2015 Simons Symposium on Tropical and Nonarchimedean Geometry.

Keywords Monge–Ampère equation • Non-Archimedean geometry • Berkovich spaces • Calabi-Yau theorem • Complex geometry • Metrics on line bundles

1 Introduction

The purpose of these notes is to discuss the Monge–Ampère equation

$$\mathrm{MA}(\phi) = \mu$$

in both the complex and non-Archimedean setting. Here μ is a positive measure¹ on the analytification of a smooth projective variety, ϕ is a semipositive metric on an ample line bundle on X , and MA is the Monge–Ampère operator. All these terms will be explained below.

In the non-Archimedean case, our presentation is based on the papers [13, 14] to which we refer for details. In the complex case, we follow [6] rather closely. Generally speaking, we avoid technicalities or detailed proofs.

¹All measures in this paper will be assumed to be Radon measures.

S. Boucksom • C. Favre
CNRS-CMLS, École Polytechnique, F-91128 Palaiseau Cedex, France
e-mail: boucksom@polytechnique.edu; charles.favre@polytechnique.edu

M. Jonsson (✉)
Department of Mathematics, University of Michigan, Ann Arbor, MI 48109-1043, USA
e-mail: mattiasj@umich.edu

2 Metrics on Lines Bundles

Let K be a field equipped with a complete multiplicative norm and let X be a smooth projective variety over K . To this data we can associate an analytification X^{an} . When K is the field of complex numbers with its usual norm, X^{an} is a compact complex manifold. When the norm is non-Archimedean, X^{an} is a K -analytic space in the sense of Berkovich [3]. In either case, it is a compact Hausdorff space.

Let L be a line bundle on X . It also admits an analytification L^{an} . A *metric* on L^{an} is a rule that to a local section $s : U \rightarrow L^{\text{an}}$, where $U \subset X^{\text{an}}$, associates a function $\|s\|$ on U , subject to the condition $\|fs\| = |f| \cdot \|s\|$, for any analytic function f on U . The metric is *continuous* if $\|s\|$ is continuous on U for every s .

For our purposes it is convenient to use additive notation for metrics and line bundles. Given an open cover U_α of X^{an} and local trivializations of L^{an} on each U_α , we can identify a section s of L with a collection $(s_\alpha)_\alpha$ of analytic functions. A metric ϕ is then a collection of functions $(\phi_\alpha)_\alpha$ in such a way that $\|s\|_\phi = |s_\alpha| e^{-\phi_\alpha}$ on U_α . With this convention, if ϕ is a metric on L^{an} , any other metric is of the form $\phi + f$, where f is a function on X^{an} . If ϕ_i is a metric on L_i , $i = 1, 2$, then $\phi_1 + \phi_2$ is a metric on $L_1 + L_2$.

Over the complex numbers, smooth metrics ϕ (i.e., each ϕ_α is smooth) play an important role. Of similar status, for K non-Archimedean, are *model metrics* defined as follows.² Let R be the valuation ring of K and k the residue field. A *model* of X is a normal scheme \mathcal{X} , flat and projective over $\text{Spec } R$ and with generic fiber isomorphic to X . A model of L is a \mathbf{Q} -line bundle \mathcal{L} on \mathcal{X} whose restriction to X is isomorphic to L . It defines a continuous metric $\phi_{\mathcal{L}}$ on L in such a way that any local nonvanishing section of a multiple of \mathcal{L} has norm constantly equal to one. Model functions, that is, model metrics on $\mathcal{O}_{\mathcal{X}}$, are dense in $C^0(X^{\text{an}})$. We refer to [21] or [14] for a more thorough discussion.

Over \mathbf{C} , a smooth metric ϕ on L^{an} is semipositive (positive) if its curvature form $dd^c\phi$ is a semipositive (positive) $(1, 1)$ -form. Here $dd^c\phi = dd^c\phi_\alpha = \frac{i}{\pi} \partial\bar{\partial}\phi_\alpha$ for any α . Such metrics only exist when L is nef.

In the non-Archimedean setting we say that a model metric $\phi_{\mathcal{L}}$ on L^{an} is semipositive if the line bundle \mathcal{L} is relatively nef, that is, its degree is nonnegative on any proper curve contained in the special fiber \mathcal{X}_0 . This implies that L is nef.

In both the complex and non-Archimedean case we say that a continuous metric ϕ is semipositive if there exists a sequence $(\phi_m)_1^\infty$ of semipositive smooth/model metrics such that $\lim_{m \rightarrow \infty} \sup_{X^{\text{an}}} |\phi_m - \phi| = 0$. In the non-Archimedean case, this notion was first introduced by Zhang [51] and Gubler [28]. In the complex case, it is more natural to say that a continuous metric ϕ is semipositive if its curvature current $dd^c\phi$ is a positive closed current. At least when L is ample, one can then

²Model metrics are not smooth in the sense of Chambert-Loir and Ducros [22] but nevertheless, for our purposes, play the same role as smooth metrics in the complex case.

prove (see Sect. 7 below) that ϕ can be approximated by smooth metrics; such an approximation is furthermore crucial for many arguments in pluripotential theory.

In the non-Archimedean case, Chambert-Loir and Ducros have introduced a notion of forms and currents on Berkovich spaces. However, it is not known whether a continuous metric whose curvature current (in their sense) is semipositive can be approximated by semipositive model metrics.

In both the complex and non-Archimedean case we denote by $\text{PSH}^0(L^{\text{an}})$ the space of continuous semipositive metrics on L^{an} . Here the superscript refers to continuity (C^0) whereas “PSH” reflects the fact that in the complex case, semipositive metrics are global versions of plurisubharmonic functions.

3 The Monge–Ampère Operator

In the complex case, the Monge–Ampère operator is a second order differential operator: we set $\text{MA}(\phi) = (dd^c \phi)^n$ for a smooth metric ϕ . It is a nonlinear operator if $n > 1$. When ϕ is semipositive, $\text{MA}(\phi)$ is a smooth positive measure on X^{an} of mass (L^n) . It is a *volume form*, that is, equivalent to Lebesgue measure, if ϕ is positive.

Next we turn to the non-Archimedean setting. *From now on we assume that K is discretely valued.* Pick a uniformizer t of the maximal ideal in the valuation ring R of K .

Consider a model metric $\phi_{\mathcal{L}}$, associated to a model $(\mathcal{X}, \mathcal{L})$ of (X, L) over $\text{Spec } R$. Write the special fiber as $\mathcal{X}_0 = \text{div}(t) = \sum_{i \in I} b_i E_i$, where E_i are the irreducible components of \mathcal{X}_0 and $b_i \in \mathbf{Z}_{>0}$. To each E_i is associated a unique (divisorial) point $x_i \in X^{\text{an}}$. We then define

$$\text{MA}(\phi) := \sum_{i \in I} b_i (\mathcal{L}|_{E_i})^n \delta_{x_i}.$$

If $\phi_{\mathcal{L}}$ is semipositive, $\mathcal{L}|_{E_i}$ is nef; hence, $(\mathcal{L}|_{E_i})^n \geq 0$ and $\text{MA}(\phi_{\mathcal{L}})$ is a positive measure. Its total mass is

$$\int_{X^{\text{an}}} 1 \cdot \text{MA}(\phi_{\mathcal{L}}) = \sum_{i \in I} b_i (\mathcal{L}|_{E_i})^n = (\mathcal{L}^n \cdot \mathcal{X}_0) = (\mathcal{L}^n \cdot \mathcal{X}_{\eta}) = (L^n).$$

Here the second to last equality follows from the flatness of \mathcal{X} over $\text{Spec } R$, and the last equality from $\mathcal{X}_{\eta} \simeq X$ and $\mathcal{L}_{\eta} \simeq L$.

From now on assume that L is *ample*, that is, we have a *polarized pair* (X, L) . In both the complex and non-Archimedean case we define $\text{MA}(\phi)$ for a continuous semipositive metric by $\text{MA}(\phi) := \lim_{m \rightarrow \infty} \text{MA}(\phi_m)$ for any sequence $(\phi_m)_1^{\infty}$ converging uniformly to ϕ . Of course, it is not obvious that the limit exists or independent of the sequence $(\phi_m)_1^{\infty}$. In the complex case this is a very special case of the Bedford–Taylor theory developed in [8, 9]. The analogous analysis in the non-Archimedean case is due to Chambert-Loir [20].

4 The Complex Monge–Ampère Equation

Theorem 4.1. *Let (X, L) be a polarized complex projective variety of dimension n and let μ be a positive measure on X^{an} of total mass (L^n) .*

- (i) *If μ is a volume form, then there exists a smooth positive metric ϕ on L^{an} such that $\text{MA}(\phi) = \mu$.*
- (ii) *If μ is absolutely continuous with respect to Lebesgue measure, with density in L^p for some $p > 1$, then there exists a (Hölder) continuous metric ϕ on L^{an} such that $\text{MA}(\phi) = \mu$.*
- (iii) *The metrics in (i) and (ii) are unique up to additive constants.*

The uniqueness statement in the setting of (i) is due to Calabi. The much harder existence part was proved by Yau [48], using PDE techniques. The combined result is often called the Calabi–Yau theorem.

The general setting of (ii) and (iii) was treated by Kołodziej [37, 38] who used methods of pluripotential theory together with a nontrivial reduction to Yau’s result. Guedj and Zeriahi [34] more generally established the existence of solutions of $\text{MA}(\phi) = \mu$ for positive measures μ (of mass (L^n)) that do not put mass on pluripolar sets. In this generality, the metrics ϕ are no longer continuous but rather lie in a suitable energy class, modeled upon work by Cegrell [19]. Dinew [24], improving upon an earlier result by Błocki [10], proved the corresponding uniqueness theorem. All these existence and uniqueness results are furthermore valid (in a suitable formulation) in the transcendental case, when (X, ω) is a Kähler manifold.

The complex Monge–Ampère equation is of fundamental importance to complex geometry. For example, it implies that every compact complex manifold with vanishing first Chern class (such manifolds are now called Calabi–Yau manifolds) admits a Ricci flat metric in any given Kähler class. The complex Monge–Ampère equation also plays a key role in recent work on the space of Kähler metrics.

5 The Non-Archimedean Monge–Ampère Equation

As before, suppose $K \simeq k((t))$ is a discretely valued field with valuation ring $R \simeq k[[t]]$ and residue field k . We further assume that K has *residue characteristic zero*, $\text{char } k = 0$. This implies that $R \simeq k[[t]]$ and $K \simeq k((t))$, where k is the residue field of K . More importantly, X then admits SNC models, that is, regular models \mathcal{X} such that the special fiber \mathcal{X}_0 has simple normal crossings. The dual complex $\Delta_{\mathcal{X}}$, encoding intersections between irreducible components of \mathcal{X}_0 , then embeds as a compact subset of X^{an} .

Theorem 5.1. *Let (X, L) be a polarized complex projective variety of dimension n over K . Assume that X is defined over a smooth k -curve. Let μ be a positive measure on X^{an} of total mass (L^n) , supported on the dual complex of some SNC model.*

- (i) *There exists a continuous metric ϕ on L^{an} such that $\text{MA}(\phi) = \mu$.*
- (ii) *The metric in (i) is unique up to an additive constant.*

Here the condition on X means that there exists a smooth projective curve C over k , a smooth projective variety Y over C , and a point $p \in C$ such that X is isomorphic to the base change $Y \times_k \text{Spec } K$, where K is the fraction field of $\widehat{\mathcal{O}}_{C,p}$. This condition is presumably redundant, but is used in the proof: see Sect. 9.

To our knowledge, the first to consider the Monge–Ampère equation (or Calabi–Yau problem) in a non-Archimedean setting were Kontsevich and Tschinkel [40]. They outlined a strategy in the case when μ is a point mass.

The case of curves ($n = 1$) was treated in detail by Thuillier in his thesis [47]; see also [2, 26]. In this case, the Monge–Ampère equation is linear and one can construct fundamental solutions by exploring the topological structure of X^{an} .

In higher dimensions, Yuan and Zhang [50] proved the uniqueness statement (ii). Their proof, based on the method by Błocki, is valid in a more general context than stated above. The first existence result was obtained by Liu [43], who treated the case when X is a maximally degenerate abelian variety and μ is equivalent to Lebesgue measure on the skeleton of X . His approach amounts to solving a *real* Monge–Ampère equation on the skeleton. The existence result (i) above was proved by the authors in [13] and the companion paper [14]. We will discuss our approach below.

The geometric ramifications of the non-Archimedean Monge–Ampère equations remain to be developed.

6 A Variational Approach

We shall present a unified approach to solving the complex and non-Archimedean Monge–Ampère equations in any dimension. The method goes back to Alexandrov’s work in convex geometry [1]. It was adapted to the complex case in [6] and to the non-Archimedean analogue in [13].

The general strategy is to construct an *energy functional*

$$E : \text{PSH}^0(L^{\text{an}}) \rightarrow \mathbf{R}$$

whose derivative is the Monge–Ampère operator, $E' = \text{MA}$, in the sense that

$$\frac{d}{dt} E(\phi + tf)|_{t=0} = \int_{X^{\text{an}}} f \text{MA}(\phi),$$

for every continuous semipositive metric $\phi \in \text{PSH}^0(L^{\text{an}})$ and every smooth/model function f on X^{an} .

Grant the existence of this functional for the moment. Given a measure μ on X^{an} , consider the functional $F_\mu : \text{PSH}^0(L^{\text{an}}) \rightarrow \mathbf{R}$ defined by

$$F_\mu(\phi) = E(\phi) - \int \phi \mu.$$

Suppose we can find $\phi \in \text{PSH}^0(L^{\text{an}})$ that maximizes F_μ . Since the derivative of F_μ is equal to $F'_\mu = \text{MA} - \mu$, we then have $0 = F'_\mu(\phi) = \text{MA}(\phi) - \mu$ as required.

Now, there are at least three problems with this approach:

- (1) There is a priori no reason why a maximizer should exist in $\text{PSH}^0(L^{\text{an}})$. We resolve this by introducing a larger space $\text{PSH}(L^{\text{an}})$ with suitable compactness properties and find a maximizer there.
- (2) Granted the existence of a maximizer $\phi \in \text{PSH}(L^{\text{an}})$, we are maximizing over a convex set rather than a vector space, so there is no reason why $F'_\mu(\phi) = 0$. Compare maximizing the function $f(x) = x^2$ on the real interval $[-1, 1]$: the maximum is not at a critical point.
- (3) In the end we want to show that—after all—the maximizer is continuous, that is, $\phi \in C^0(L^{\text{an}})$.

We shall discuss how to address (1) and (2) in the next two sections. The continuity result in (3) requires a priori capacity estimates due to Kołodziej, and will not be discussed in these notes.

7 Singular Semipositive Metrics

Plurisubharmonic (psh) functions are among the *objets souples* (soft objects) in complex analysis according to P. Lelong [42]. This is reflected in certain useful compactness properties. The global analogues of psh functions are semipositive singular metrics on holomorphic line bundles. Here “singular” means that vectors may have infinite length.

Theorem 7.1. *Let K be either \mathbf{C} or a discretely valued field of residue characteristic zero, and let (X, L) be a smooth projective polarized variety over K . Then there exists a unique class $\text{PSH}(L^{\text{an}})$, the set of singular semipositive metrics, with the following properties:*

- $\text{PSH}(L^{\text{an}})$ is a convex set which is closed under maxima and addition of constants;
- $\text{PSH}(L^{\text{an}}) \cap C^0(L^{\text{an}}) = \text{PSH}^0(L^{\text{an}})$;
- if s_i , $1 \leq i \leq p$, are nonzero global sections of mL for some $m \geq 1$, then $\phi := \frac{1}{m} \max_i \log |s_i| \in \text{PSH}(L^{\text{an}})$; further, ϕ is continuous iff the sections s_i have no common zero;

- if (ϕ_j) is an arbitrary family in $\text{PSH}(L^{\text{an}})$ that is uniformly bounded from above, then the usc regularization of $\sup_j \phi_j$ belongs to $\text{PSH}(L^{\text{an}})$;
- if (ϕ_j) is a decreasing net in $\text{PSH}(L^{\text{an}})$, then either $\phi_j \rightarrow -\infty$ uniformly on X^{an} or $\phi_j \rightarrow \phi$ pointwise on X^{an} for some $\phi \in \text{PSH}(L^{\text{an}})$;
- **Regularization:** for every $\phi \in \text{PSH}(L^{\text{an}})$ there exists a decreasing sequence $(\phi_m)_{m=1}^\infty$ of smooth/model metrics such that ϕ_m converges pointwise to ϕ on X^{an} as $m \rightarrow \infty$; and
- **Compactness:** the space $\text{PSH}(L^{\text{an}})/\mathbf{R}$ is compact.

To make sense of the compactness statement we need to specify the topology on $\text{PSH}(L^{\text{an}})$. In the complex case, one usually fixes a volume form μ on X^{an} and takes the topology induced by the L^1 -norm: $\|\phi - \psi\| = \int_{X^{\text{an}}} |\phi - \psi| \mu$. In the non-Archimedean case, there is typically no volume form on X^{an} . Instead, we say that a net $(\phi_j)_j$ in $\text{PSH}(L^{\text{an}})$ converges to ϕ if $\lim_j \sup_{\Delta_{\mathcal{X}}} |\phi_j - \phi| = 0$ for every SNC model \mathcal{X} . Implicit in this definition is that the restriction to $\Delta_{\mathcal{X}}$ of every singular metric in $\text{PSH}(L^{\text{an}})$ is continuous: see Theorem 7.2 below.

In the complex case, one typically defines $\text{PSH}(L^{\text{an}})$ as the set of usc singular metrics ϕ that are locally represented by L^1 functions and whose curvature current $dd^c \phi$ (computed in the sense of distributions) is a positive closed current. Thus ϕ is locally given as the sum of a smooth function and a psh function. Most of the statements above then follow from basic facts about plurisubharmonic functions in \mathbf{C}^n . The regularization result is the most difficult. On \mathbf{C}^n it is easy to regularize using convolutions. With some care, one can in the global (projective) case glue together local regularizations to obtain a global one. See [23] for a general result and [11] for a relatively simple argument applicable in our setting.

In the non-Archimedean case, we are not aware of any workable a priori definition of $\text{PSH}(L^{\text{an}})$. Chambert-Loir and Ducros [22] have a notion of forms and currents on Berkovich spaces, but it is unclear if it gives the right objects for the purposes of the theorem above. Instead, we prove the following result:

Theorem 7.2. *For any SNC model \mathcal{X} , the restriction of the dual complex $\Delta_{\mathcal{X}} \subset X^{\text{an}}$ of the set of model metrics on L^{an} forms an equicontinuous family.*

This is proved using a rather subtle argument, involving intersection numbers on toroidal models dominating \mathcal{X} . It would be interesting to have a different proof. At any rate, Theorem 7.2 allows us to define $\text{PSH}(L^{\text{an}})$ as the set of usc singular metrics ϕ satisfying, for every sufficiently large SNC model \mathcal{X} ,

- (i) $(\phi - \phi_0) \circ r_{\mathcal{X}} \geq \phi - \phi_0$ and
- (ii) the restriction of ϕ to $\Delta_{\mathcal{X}}$ is a uniform limit of a sequence $\phi_m|_{\Delta_{\mathcal{X}}}$, where each ϕ_m is a semipositive model metric.

Here ϕ_0 is a fixed model metric, determined by some model dominated by \mathcal{X} . The map $r_{\mathcal{X}} : X^{\text{an}} \rightarrow \Delta_{\mathcal{X}} \subset X^{\text{an}}$ is a natural retraction. Since ϕ is usc, condition (i) implies that $\phi = \phi_0 + \lim_{\mathcal{X}} (\phi - \phi_0) \circ r_{\mathcal{X}}$, so that ϕ is determined by its restrictions to all dual complexes.

With this definition, the compactness of $\text{PSH}(L^{\text{an}})/\mathbf{R}$ follows from Theorem 7.2 and Ascoli's theorem. Regularization, however, is quite difficult to show. We are not aware of any procedure that would replace convolution in the complex case. Instead we use algebraic geometry. Here is an outline of the proof.

Fix $\phi \in \text{PSH}(L^{\text{an}})$. For any SNC model \mathcal{X} , ϕ naturally induces a model metric $\phi_{\mathcal{X}}$. The semipositivity of ϕ implies that the net $(\phi_{\mathcal{X}})_{\mathcal{X}}$ indexed by the collection of (isomorphism classes of) SNC models decreases to ϕ . Unfortunately, except in the curve case $n = 1$, $\phi_{\mathcal{X}}$ has no reason to be semipositive; this reflects the fact that the pushforward of a nef line bundle may fail to be nef. We address this by defining $\psi_{\mathcal{X}}$ as the supremum of all semipositive (singular) metrics dominated by $\phi_{\mathcal{X}}$. We then show that $\psi_{\mathcal{X}}$ is continuous and can be *uniformly* approximated by a sequence $(\psi_{\mathcal{X},m})_m^{\infty}$ of semipositive model metrics. From this data it is not hard to produce a decreasing net of semipositive model metrics converging to ϕ .

Let us say a few words on the construction of the semipositive model metrics $\phi_{\mathcal{X},m}$ since this is a key step in the paper [14]. For simplicity assume that L is base point free and that $\phi_{\mathcal{X}}$ is associated with a line bundle \mathcal{L} (rather than an \mathbf{R} -line bundle) on \mathcal{X} . Let \mathfrak{a}_m be the base ideal of $m\mathcal{L}$, cut out by the global sections; it is cosupported on the special fiber \mathcal{X}_0 . The sequence $(\mathfrak{a}_m)_m$ is a *graded sequence* in the sense that $\mathfrak{a}_l \cdot \mathfrak{a}_m \subset \mathfrak{a}_{l+m}$. Each \mathfrak{a}_m naturally defines a semipositive model metric $\psi_{\mathcal{X},m}$ on L^{an} . The fact that $\psi_{\mathcal{X},m}$ converges *uniformly* to $\psi_{\mathcal{X}}$ translates into a statement that the graded sequence $(\mathfrak{a}_m)_m$ is “almost” finitely generated. This in turn is proved using *multiplier ideals* and ultimately reduces to the Kodaira vanishing theorem; to apply the latter, it is crucial to work in residue characteristic zero.

The argument above proves that any $\phi \in \text{PSH}(L^{\text{an}})$ is the limit of a decreasing *net* of semipositive model metrics. When ϕ is continuous, the convergence is uniform by Dini's theorem, and we can use the sup-norm to extract a decreasing *sequence* of model metrics converging to ϕ . In the general case, the *Monge–Ampère capacity* developed in [13, § 4] (and modeled on [8, 33]) can similarly be used to extract a convergent sequence from a net.

8 Energy

In the complex case, the (Aubin–Mabuchi) energy functional is defined as follows. Fix a smooth semipositive reference metric ϕ_0 and set

$$E(\phi) := \frac{1}{n+1} \sum_{j=0}^n \int_{X^{\text{an}}} (\phi - \phi_0)(dd^c \phi)^j \wedge (dd^c \phi_0)^{n-j}. \quad (1)$$

for any smooth metric ϕ . Here $(dd^c \phi)^j \wedge (dd^c \phi_0)^{n-j}$ is a *mixed Monge–Ampère measure*. It is a positive measure if ϕ is semipositive.

In the non-Archimedean case, mixed Monge–Ampère measures can be defined using intersection theory when ϕ and ϕ_0 are model metrics, and the energy of ϕ is then defined exactly as above.

For two smooth/model metrics ϕ and ψ we have

$$E(\phi) - E(\psi) = \frac{1}{n+1} \sum_{j=0}^n \int_{X^{\text{an}}} (\phi - \psi)(dd^c \phi)^j \wedge (dd^c \psi)^{n-j}. \tag{2}$$

This is proved using integration by parts in the complex case and follows from basic intersection theory in the non-Archimedean case.

We can draw two main conclusions from (2). First, the derivative of the energy functional is the Monge–Ampère operator, in the sense that

$$\left. \frac{d}{dt} E(\phi + tf) \right|_{t=0} = \int_{X^{\text{an}}} f \text{MA}(\phi) \tag{3}$$

for a smooth/model metric ϕ on L^{an} and a smooth/model function f on X^{an} .

Second, $E(\psi) \geq E(\phi)$ when $\psi \geq \phi$ are semipositive. It then makes sense to set

$$E(\phi) := \inf\{E(\psi) \mid \psi \geq \phi, \psi \text{ a semipositive smooth/model metric on } L^{\text{an}}\}.$$

for any singular semipositive metric $\phi \in \text{PSH}(L^{\text{an}})$. The resulting functional

$$E : \text{PSH}(L^{\text{an}}) \rightarrow [-\infty, \infty)$$

has many good properties: E is concave, monotonous, and satisfies $E(\phi + c) = E(\phi) + c$ for $c \in \mathbf{R}$. Further, E is usc and continuous along decreasing nets.

The energy functional singles out a class $\mathcal{E}^1(L^{\text{an}})$ of metrics with *finite energy*, $E(\phi) > -\infty$. This class has good properties. In particular, one can (with some effort) define mixed Monge–Ampère measures $(dd^c \phi)^j \wedge (dd^c \psi)^{n-j}$ for $\phi, \psi \in \mathcal{E}^1(L^{\text{an}})$, and (1) continues to hold.

Let us now go back to the variational approach to solving the Monge–Ampère equation. Fix a positive measure μ on X^{an} of mass (L^n) . In the complex case we assume that μ is absolutely continuous with respect to Lebesgue measure, with density in L^p for some $p > 1$. In the non-Archimedean case we assume that μ is supported on some dual complex. In both cases, one can show that the functional $\phi \rightarrow \int (\phi - \phi_0) \mu$ is (finite and) continuous on $\text{PSH}(L^{\text{an}})$, where ϕ_0 is the same reference metric as in (1). Thus the functional $F_\mu : \text{PSH}(L^{\text{an}}) \rightarrow [-\infty, \infty)$ defined by

$$F_\mu(\phi) := E(\phi) - \int (\phi - \phi_0) \mu$$

is upper semicontinuous. It follows from (2) that F_μ does not depend on the choice of reference metric ϕ_0 . We also have $F_\mu(\phi + c) = F_\mu(\phi)$ for $\phi \in \text{PSH}(L^{\text{an}})$, $c \in \mathbf{R}$. Thus F_μ descends to a usc functional on the quotient space $\text{PSH}(L^{\text{an}})/\mathbf{R}$. By Theorem 7.1, the latter space is compact, so we can find $\phi \in \text{PSH}(L^{\text{an}})$ maximizing F_μ . It is clear that $\phi \in \mathcal{E}^1(L^{\text{an}})$, so the mixed Monge–Ampère measures of ϕ and ϕ_0 are well defined. However, equation (3) no longer makes sense, since there is no reason for the metric $\phi + tf$ to be semipositive for $t \neq 0$. Therefore, it is not clear that $\text{MA}(\phi) = \mu$, as desired. In the next section, we explain how to get around this problem.

9 Envelopes, Differentiability, and Orthogonality

We define the *psh envelope* of a (possibly singular) metric ψ on L^{an} by

$$P(\psi) := \sup\{\phi \in \text{PSH}(L^{\text{an}}) \mid \phi \leq \psi\}^*.$$

As before, ϕ^* denotes the usc regularization of a singular metric ϕ . In all cases, we need to consider, ψ will be the sum of a metric in $\mathcal{E}^1(L^{\text{an}})$ and a continuous function on X^{an} . In particular, ψ is usc, $P(\psi) \in \mathcal{E}^1(L^{\text{an}})$ and $P(\psi) \leq \psi$.

This envelope construction was in fact already mentioned at the end of Sect. 7 as it plays a key role in the regularization theorem. The psh envelope is an analogue of the convex hull; see Fig. 1.

The key fact about the psh envelope is that the composition $E \circ P$ is differentiable and that $(E \circ P)' = E' \circ P$. More precisely, we have

Theorem 9.1. *For any $\phi \in \mathcal{E}^1(L^{\text{an}})$ and $f \in C^0(X^{\text{an}})$, the function $t \mapsto E(P(\phi + tf))$ is differentiable at $t = 0$, with derivative $\frac{d}{dt}E(\phi + tf)|_{t=0} = \int f \text{MA}(\phi)$.*


Granted this result, let us show how to solve the Monge–Ampère equation. Pick $\phi \in \mathcal{E}^1(L^{\text{an}})$ that maximizes $F_\mu(\phi) = E(\phi) - \int (\phi - \phi_0)\mu$ and consider any $f \in C^0(X^{\text{an}})$. For any $t \in \mathbf{R}$ we have

$$\begin{aligned} E(P(\phi + tf)) - \int (\phi + tf - \phi_0)\mu &\leq E(P(\phi + tf)) - \int (P(\phi + tf) - \phi_0)\mu \\ &\leq E(\phi) - \int (\phi - \phi_0)\mu. \end{aligned}$$

Since the left-hand side is differentiable at $t = 0$, the derivative must be zero, which amounts to $\int f \text{MA}(\phi) - \int f \mu = 0$. Since $f \in C^0(X^{\text{an}})$ was arbitrary, this means that $\text{MA}(\phi) = \mu$, as desired.

The proof of this differentiability results proceeds by first reducing to the case when ϕ and f are continuous. A key ingredient is then

Fig. 1 The convex hull $P(f)$ of a continuous function f of one variable. Note that $P(f)$ is affine, i.e., $P(f)'' = 0$ where $P(f) \neq f$



Theorem 9.2. For any continuous metric ϕ on L^{an} we have

$$\int_{X^{\text{an}}} (\phi - P(\phi)) \text{MA}(P(\phi)) = 0. \tag{4}$$

In other words, the Monge–Ampère measure $\text{MA}(P(\phi))$ is supported on the locus $P(\phi) = \phi$. A version of this for functions of one variable is illustrated in Fig. 1.

To prove this result, we can reduce to the case when ϕ is a smooth/model metric. In the complex case, Theorem 9.2 was proved by Berman and the first author in [5] using the pluripotential theoretic technique known as “balayage.” In the non-Archimedean setting, Theorem 9.2 is deduced in [13] from the asymptotic orthogonality of Zariski decompositions in [12] and is for this reason called the *orthogonality property*. The assumption in Theorem 5.1 that the variety X be defined over a smooth k -curve is used exactly in order to apply the result from [12].

The solution to the non-Archimedean Monge–Ampère equation $\text{MA}(\phi) = \mu$ can be made slightly more explicit in the case when the support of μ is a singleton, $\mu = d_L \delta_x$, where $d_L := (L^n)$ and $x \in X^{\text{an}}$ belongs to some dual complex; such points x are known as *quasimonomial* or *Abhyankar* points.

The fiber L_x^{an} of L^{an} above $x \in X^{\text{an}}$ is isomorphic to the Berkovich affine line over the complete residue field $\mathcal{H}(x)$. Fix any nonzero $y \in L_x^{\text{an}}$ and set

$$\phi_x := \sup\{\phi \in \text{PSH}(L^{\text{an}}) \mid \|y\|_{\phi} \geq 1\}.$$

By Boucksom et al. [13, Proposition 8.6], $\text{MA}(\phi_x)$ is supported on x , so $\text{MA}(\phi_x) = d_L \delta_x$. It would be interesting to find an example of a divisorial point $x \in X^{\text{an}}$ such that ϕ_x is not a model function.

10 Curves

The disadvantage of the variational approach to the Monge–Ampère equation is that it gives very little control on the solution beyond continuity. Here we shall make the solution more concrete in the case of curves; the next section deals with toric varieties.

Thus assume that X is a smooth projective curve over K . In this case, the Monge–Ampère operator (which one would normally refer to as the Laplacian) is *linear*: if ϕ_i is a metric on $L_i, i = 1, 2$, then $\text{MA}(\phi_1 + \phi_2) = \text{MA}(\phi_1) + \text{MA}(\phi_2)$. Furthermore,

as we shall see, the Monge–Ampère operator is naturally defined on *any* singular semipositive metric on L^{an} , for an ample line bundle L , and we can solve $\text{MA}(\phi) = \mu$ for *any* positive measure μ of mass $\text{deg } L$.

Let us first explain this in the complex case; X^{an} is then a compact Riemann surface. Fix a smooth metric ϕ_0 on L^{an} . The curvature form $\omega_0 := dd^c \phi_0$ is a volume form of mass $\text{deg } L$. A singular metric ϕ on L^{an} is then semipositive iff $\varphi := \phi - \phi_0$ is an ω_0 -psh function, that is, a locally integrable function φ that is locally the sum of a smooth function and a psh function, and such that $\omega_0 + dd^c \varphi$ is a positive measure. We then set $\text{MA}(\phi) := \omega_0 + dd^c \varphi$; this definition does not depend on the choice of ϕ_0 .

Now we explain how to solve the equation $\text{MA}(\phi) = \mu$ for any positive measure μ of mass $d_L := \text{deg } L$. Writing $\phi = \phi_0 + \varphi$ as above, we must solve $dd^c \varphi = \mu - \omega_0$, where $\omega_0 := dd^c \phi_0$. It suffices to do this when $\mu = d_L \delta_x$ for some $x \in X^{\text{an}}$: indeed, if we normalize the solution φ_x to $dd^c \varphi_x = \mu - d_L \delta_x$ by $\int_{X^{\text{an}}} \varphi_x \omega_0 = 0$, then the function φ_μ defined by $\varphi_\mu(y) := d_L^{-1} \int_{X^{\text{an}}} \varphi_x(y) d\mu(x)$ satisfies $dd^c \varphi_\mu = \omega_0 - \mu$ and is normalized by $\int_{X^{\text{an}}} \varphi_\mu \omega_0 = 0$.

The function φ_x can be “physically” interpreted as the voltage (suitably normalized) when putting a charge of $+d_L$ at the point x and a total charge of $-d_L$ spread out according to the measure ω_0 . Mathematically, Perron’s method describes it as the supremum of all ω_0 -subharmonic functions φ on X^{an} satisfying $\int_{X^{\text{an}}} \varphi \omega_0 = 0$ and $\varphi \leq d_L \log |z| + O(1)$, where z is a local coordinate at x .

Now we consider the non-Archimedean case. As before, let us assume that K is a discretely valued field of residue characteristic zero, even though this is not really necessary in the one-dimensional case.³

The main point is that any Berkovich curve has the structure of a generalized⁴ metric graph. We will not describe this in detail, but here is the idea. The dual graph $\Delta_{\mathcal{X}}$ of any SNC model \mathcal{X} is a connected, one-dimensional simplicial complex. As before, we view it as a subset of X^{an} . It carries a natural integral affine structure, inducing a metric. If \mathcal{X}' is an SNC model dominating \mathcal{X} (in the sense that the canonical birational map $\mathcal{X} \dashrightarrow \mathcal{X}'$ is a morphism), then $\Delta_{\mathcal{X}}$ is a subset of $\Delta_{\mathcal{X}'}$ and the inclusion $\Delta_{\mathcal{X}} \rightarrow \Delta_{\mathcal{X}'}$ is an isometry. There is also a (deformation) retraction $r_{\mathcal{X}} : X^{\text{an}} \rightarrow \Delta_{\mathcal{X}}$, and $X^{\text{an}} \simeq \varprojlim_{\mathcal{X}} \Delta_{\mathcal{X}}$. In this way, the metrics on the dual complexes induce a generalized metric on X^{an} .

The structure of each $\Delta_{\mathcal{X}}$ and of X^{an} as metric graphs allows us to define a Laplacian on these spaces, by combining the real Laplacian on segments and the combinatorial Laplacian at branch points (and endpoints). This Laplacian allows us to understand both semipositive singular metrics and the Monge–Ampère operator.

Namely, fix a model metric ϕ_0 on L^{an} . It is represented by a \mathbf{Q} -line bundle on some SNC model $\mathcal{X}^{(0)}$. The measure $\omega_0 := dd^c \phi_0$ is supported on the vertices of $\Delta_{\mathcal{X}^{(0)}}$. Now, a singular metric ϕ is semipositive iff for every SNC model \mathcal{X} dominating $\mathcal{X}^{(0)}$, the restriction of the function $\phi - \phi_0$ to $\Delta_{\mathcal{X}}$ is a ω_0 -subharmonic

³Indeed, Thuillier [47] systematically develops a potential theory on Berkovich curves in a very general setting.

⁴This means that some distances may be infinite.

function in the sense that $\Delta((\phi - \phi_0)|_{\Delta_{\mathcal{X}}}) = \mu_{\mathcal{X}} - \omega_0$, where $\mu_{\mathcal{X}}$ is a positive measure on $\Delta_{\mathcal{X}}$ of mass d_L . In this case, there further exists a unique measure μ on X^{an} of mass d_L such that $(r_{\mathcal{X}})_*\mu = \mu_{\mathcal{X}}$ for all \mathcal{X} , and we have $\text{MA}(\phi) = \mu$.

To solve the equation $\text{MA}(\phi) = \mu$ for a positive measure μ of mass d_L , it suffices by linearity to treat the case $\mu = d_L\delta_x$ for a point $x \in X^{\text{an}}$. In this case, the function $\varphi := \phi - \phi_0$ will be locally constant outside the convex hull of $\{x\} \cup \Delta_{\mathcal{X}^{(0)}}$. The latter is essentially a finite metric graph on which we need to find a function whose Laplacian is equal to $d_L\delta_x - \omega_0$. This can be done in a quite elementary way.

An interesting example of semipositive metrics, both in the complex and non-Archimedean case, comes from dynamics [51]. Suppose $f : (X, L) \curvearrowright$ is a polarized endomorphism of degree $\lambda > 1$. In other words, $f : X \rightarrow X$ is an endomorphism and f^*L is linearly equivalent to λL . Then there exists a unique *canonical metric* ϕ_{can} on L^{an} , satisfying $f^*\phi = \lambda\phi$. This metric is continuous and semipositive but usually not a model metric.

As a special case, suppose X is an elliptic curve and that f is the map given by multiplication by λ . In the complex case, $X^{\text{an}} \simeq \mathbf{C}/\Lambda$ is a torus and $\mu_{\text{can}} := \text{MA}(\phi_{\text{can}})$ is given by a multiple of Haar measure on X^{an} . In the non-Archimedean case, there are two possibilities. If X has good reduction over $\text{Spec } R$, then μ_{can} is a point mass. Otherwise, μ_{can} is proportional to Lebesgue measure on the *skeleton* $\text{Sk}(X^{\text{an}})$, a subset homeomorphic to a circle. A similar description of the measure μ_{can} in the case of higher-dimensional abelian varieties is given in [29].

11 Toric Varieties

For general facts about toric varieties, see [18, 27, 36]. In this section we briefly describe how the complex and non-Archimedean points of view elegantly come together in the toric setting and translate into statements about convex functions and the real Monge–Ampère operator. As before, we only consider the non-Archimedean field $K = k((t))$ with $\text{char } k = 0$; however, most of what we say here should be true in a more general context: see [30].

Let $M \simeq \mathbf{Z}^n$ be a free abelian group, N its dual, and let $T = \text{Spec } K[M]$ be the corresponding split K -torus. A polarized toric variety (X, L) is then determined by a rational polytope $\Delta \subset M_{\mathbf{R}}$. The variety X is described by the normal fan to Δ in $N_{\mathbf{R}}$ and the points of $M \cap \Delta$ are in 1–1 correspondence with equivariant sections of L ; we write χ^u for the section of L associated with $u \in M$. This description is completely general and holds over any field as well as over \mathbf{Z} .

There is also a “tropical” space X^{trop} associated with X . As a topological space, it is compact and contains $N_{\mathbf{R}}$ as an open dense subset.⁵ For any valued field K , there

⁵In our setting, X^{trop} can be identified with the (moment) polytope Δ in such a way that $N_{\mathbf{R}}$ corresponds to the interior of Δ , but this identification does not preserve the affine structure on $N_{\mathbf{R}}$.

is a tropicalization map $\text{trop} : X^{\text{an}} \rightarrow X^{\text{trop}}$, where X^{an} refers to the analytification with respect to the norm on K . The inverse image of $N_{\mathbf{R}}$ is the torus T^{an} .

There is a natural correspondence between equivariant metrics on L^{an} and functions on $N_{\mathbf{R}}$. Let ϕ is an equivariant metric on L^{an} . For every $u \in M$, χ^u is a nonvanishing section of L on T so $\phi - \log |\chi^u|$ defines a function on T^{an} that is constant on the fibers of the tropicalization map. In particular, picking $u = 0$, we can write

$$\phi - \log |\chi^0| = g \circ \text{trop} \tag{5}$$

for some function g on $N_{\mathbf{R}}$. Conversely, given a function g on $N_{\mathbf{R}}$, (5) defines an equivariant metric on the restriction of L^{an} to T^{an} .

We now go from the torus T to the polarized variety (X, L) . After replacing L by a multiple, we may assume that all the vertices of Δ belong to M . Set

$$\phi_{\Delta} := \max_{u \in \Delta} \log |\chi^u|.$$

This is a semipositive, equivariant model metric on L^{an} . Its restriction to T^{an} corresponds to the homogeneous, nonnegative, convex function

$$g_{\Delta} := \max_{u \in \Delta} u$$

on $N_{\mathbf{R}}$. In general, an equivariant singular metric ϕ on L^{an} corresponds to a convex function g on $N_{\mathbf{R}}$ such that $g \leq g_{\Delta} + O(1)$. It is bounded iff $g - g_{\Delta}$ is bounded on $N_{\mathbf{R}}$.

The real Monge–Ampère measure of any convex function g on $N_{\mathbf{R}}$ is a well-defined positive measure $\text{MA}_{\mathbf{R}}(g)$ on $N_{\mathbf{R}}$ (see, e.g., [46]). When $g = g_{\Delta} + O(1)$, its total mass is given by

$$\int_{N_{\mathbf{R}}} \text{MA}_{\mathbf{R}}(g) = \text{Vol}(\Delta) = \frac{(L^n)}{n!},$$

where the last equality follows from [27, p. 111].

We now wish to relate the real Monge–Ampère measure of g and the Monge–Ampère measure of the corresponding semipositive metric ϕ on L^{an} .

First consider the non-Archimedean case, in which there is a natural embedding $j : N_{\mathbf{R}} \rightarrow T^{\text{an}} \subset X^{\text{an}}$ given by monomial valuations that sends $v \in N_{\mathbf{R}}$ to the norm

$$\sum_{u \in M} a_u u \in K[M] \mapsto \max_{u \in M} \{|a_u| \exp(-\langle u, v \rangle)\}.$$

In particular, $j(0) = x_G$, the Gauss point of the open T -orbit.

If g is a convex function on $N_{\mathbf{R}}$ with $g = g_{\Delta} + O(1)$, and if ϕ is the corresponding continuous semipositive metric on L , then [18, Theorem 4.7.4] asserts that

$$\text{MA}(\phi) = n! j_* \text{MA}_{\mathbf{R}}(g).$$

For a compactly supported positive measure ν on $N_{\mathbf{R}}$ of mass (L^n) , solving the Monge–Ampère equation $\text{MA}(\phi) = j_*(\nu)$ therefore amounts to solving the real Monge–Ampère equation $\text{MA}_{\mathbf{R}}(g) = \nu/n!$. This can be done explicitly when ν is a point mass, say, supported at $v_0 \in N_{\mathbf{R}}$. Indeed, the function $g_{v_0} : N \rightarrow \mathbf{R}$ defined by $g = g_{\Delta}(-v_0)$ is convex and satisfies $g = g_{\Delta} + O(1)$. Further, for every point $v \neq v_0$ there exists a line segment in $N_{\mathbf{R}}$ containing v in its interior and on which g is affine. This implies that $\text{MA}_{\mathbf{R}}(g)$ is supported at v_0 . As a consequence, the corresponding continuous metric ϕ on L^{an} satisfies $\text{MA}_{\mathbf{R}}(\phi) = (L^n)\delta_{j(u_0)}$.

This solution can be shown to tie in well with the construction at the end of Sect. 9, but is of course much more explicit. For example, when $u_0 \in N_{\mathbf{Q}}$, so that $j(u_0) \in X^{\text{an}}$ is divisorial, the function g_{u_0} is \mathbf{Q} -piecewise linear so that the corresponding metric ϕ is a model metric.

Finally we consider the complex case. In this case we cannot embed $N_{\mathbf{R}}$ in T^{an} . However, the preimage of any point $v \in N_{\mathbf{R}}$ under the tropicalization is a real torus of dimension n in T^{an} on which the multiplicative group $(S^1)^n$ acts transitively. To any compactly supported positive measure ν on $N_{\mathbf{R}}$ of mass $(L^n)/n!$ we can therefore associate a unique measure μ on T^{an} , still denoted $\mu := j_*\nu$, that is invariant under the action of $(S^1)^n$ and satisfies $\text{trop}_*\mu = \nu$.

If ϕ is an equivariant semipositive metric on L^{an} , corresponding to a convex function g on $N_{\mathbf{R}}$, we then have

$$\text{MA}(\phi) = n!j_*\text{MA}_{\mathbf{R}}(g).$$

For $(S^1)^n$ -invariant measures μ on L^{an} of mass (L^n) , solving the complex Monge–Ampère equation $\text{MA}(\phi) = \mu$ thus reduces to solving the real Monge–Ampère equation $\text{MA}_{\mathbf{R}}(g) = \frac{1}{n!}\text{trop}_*\mu$.

12 Outlook

In this final section we indicate some possible extensions of our work and make a few general remarks.

First of all, it would be nice to have a *local* theory for semipositive singular metrics. Indeed, while the global approach in [13, 14] works well for the Calabi–Yau problem, it has some unsatisfactory features. For example, it is not completely trivial to prove that the Monge–Ampère operator is local in the sense that if ϕ_1, ϕ_2 are two (say) continuous semipositive metrics that agree on an open subset $U \subset X^{\text{an}}$, then $\text{MA}(\phi_1) = \text{MA}(\phi_2)$ on U . We prove this in [13] using the Monge–Ampère capacity. Still, it would be desirable to say that the restriction of a semipositive metric to (say) an open subset of X^{an} remains semipositive!

In contrast, in the complex case, the classical approach is local in nature. Namely, one first defines and studies psh functions on open subsets of \mathbf{C}^n and then defines

singular semipositive metrics as global analogues. By construction, the Monge–Ampère operator is a local (differential) operator.⁶

In a general non-Archimedean setting, Chambert-Loir and Ducros [22] (see also [31, 32]) define psh functions as continuous functions φ such that $d'd''\varphi$ is a positive closed current (in their sense), for suitable operators d' and d'' analogous to their complex counterparts and modeled on notions due to Lagerberg [41]. While this leads to a very nice theory, that moreover works for general Berkovich spaces, the crucial compactness and regularization results are so far missing. At any rate, the tropical charts used in [22] may be a good substitute for dual complexes of SNC models.

Going back to the projective setting, there are several open questions and possible extensions, even in the case of a discretely valued ground field of residue characteristic zero.

First, when solving the Monge–Ampère equation, we needed to assume that the variety X was obtained by base change from a variety over a k -curve. This assumption was made in order to use the orthogonality result in [12], but is presumably redundant.

Second, one should be able to solve the Monge–Ampère equation $\text{MA}(\phi) = \mu$ for more general measures μ . In the complex setting, this is done in [24, 34] for non-pluripolar measures μ . The analogous result should be valid in the non-Archimedean setting, too, although some countability issues seem to require careful attention. Having such a general result would allow for a nice Legendre duality, as explored in [4, 6] in the complex case.

Third, one could try to get more specific information about the solution. We already mentioned at the end of Sect. 9 that we don't know whether the solution to the equation $\text{MA}(\phi) = d_L \mu_x$ is a model function for x a divisorial point (and $d_L = (L^n)$). In a different direction, one could consider the case when X is a Calabi–Yau variety, in the sense that $K_X \simeq \mathcal{O}_X$. Then there exists a canonical subset $\text{Sk}(X) \subset X^{\text{an}}$, the *Kontsevich–Soibelman skeleton*, see [39, 44, 45]. It is a subcomplex of the dual complex of any SNC model and comes equipped with an integral affine structure, inducing a volume form on each face. One can solve $\text{MA}(\phi) = \mu$, for linear combinations of these volume forms, viewed as measures on X^{an} . Can we say anything concrete about the solution ϕ , as in the case of maximally degenerate abelian varieties considered in [43]?

It would obviously be interesting to work over other types of non-Archimedean fields, such as \mathbf{Q}_p . Here there are several challenges. First, we systematically use SNC models, which are only known to exist in residue characteristic zero (except in low dimensions). It is possible that the tool of SNC models can, with some additional effort, be replaced by alterations, tropical charts, or other methods. However, we also crucially use the assumption of residue characteristic

⁶However, one also needs to verify that the Monge–Ampère operator is local for the *plurifine* topology. This is nontrivial in both the complex and non-Archimedean case.

zero when applying the vanishing theorems that underlie the regularization theorem for singular semipositive metrics. Here some new ideas are needed.

A simpler situation to handle is that of a *trivially* valued field. This is explored in [17] and can be briefly explained as follows. Let k be any field of characteristic zero, equipped with the trivial norm. Let (X, L) be a polarized variety over k . In this setting, the notion of model metrics and model functions seemingly does not take us very far, as the only model of X is X itself! Instead, the idea is to use a non-Archimedean field extension. Set $K = k((t))$, $X_K := X \otimes_k K$, etc. The multiplicative group $G := \mathbf{G}_{m,k}$ acts on X_K^{an} and X^{an} can be identified with the set of G -equivariant points in X_K^{an} . Similarly, singular semipositive metrics on L^{an} are defined as G -invariant singular semipositive metrics on L_K^{an} . In this way, the main results about $\text{PSH}(L^{\text{an}})$ follow from the corresponding results about $\text{PSH}(L_K^{\text{an}})$ and the same is true for the solution of the Monge–Ampère equation.

A primary motivation for studying the trivially valued case, at least in the case $k = \mathbf{C}$, is that the space of singular semipositive metrics on L^{an} naturally sits “at the boundary” of the space of positive (Kähler) metrics on the holomorphic line bundle L . As such, it can be used to study questions on K -stability and may be useful for the study of the existence of constant scalar curvature metrics, see [15, 16]. A different scenario where a complex situation degenerates to a non-Archimedean one occurs in [35].

In yet another direction, one could try to consider line bundles that are not necessarily ample, but rather big and nef, or simply big. In the complex case this was done in [7, 25]. One motivation for such a generalization is that it is invariant under birational maps and would hence allow us to study singular varieties.

Finally, it would be interesting to have transcendental analogues. Indeed, in the complex case, one often starts with a Kähler manifold X together with a Kähler class ω , rather than a polarized pair (X, ω) . A notion of Kähler metric is proposed in [40, 49], but it is not clear whether or not this plays the role of a (possibly) transcendental Kähler metric.

Acknowledgements These notes are partially based on talks given at the 2015 Simons Symposium on Tropical and Nonarchimedean Geometry by the first and third authors. We want to thank the Simons foundation as well as the organizers and participants of the symposium for the opportunity to present our work, and for stimulating discussions. Boucksom was supported by the ANR projects MACK and POSITIVE. Favre was supported by the ERC-starting grant project “INonarcomp” I no.307856. Jonsson was supported by NSF grant DMS-1266207.

References

1. Alexandrov, A.D.: On the theory of mixed volumes of convex bodies III: Extension of two theorems of Minkowski on convex polyhedra to arbitrary convex bodies. (Russian) *Mat. Sbornik* **3**(1), 27–44 (1938) [English translation available in *Selected works. Part I. Selected scientific papers*. Gordon and Breach, Amsterdam, 1996]
2. Baker, M., Rumely, R.: *Potential Theory on the Berkovich Projective Line*. *Mathematical Surveys and Monographs*, vol. 159. American Mathematical Society, Providence, RI (2010)

3. Berkovich, V.G.: Spectral theory and analytic geometry over non-Archimedean fields. *Mathematical Surveys and Monographs*, vol. 33. American Mathematical Society, Providence, RI (1990)
4. Berman, R.: A thermodynamical formalism for Monge–Ampère equations, Moser–Trudinger inequalities and Kähler–Einstein metrics. *Adv. Math.* **248**, 1254–1297 (2013)
5. Berman, R., Boucksom, S.: Growth of balls of holomorphic sections and energy at equilibrium. *Invent. Math.* **181**, 337–394 (2010)
6. Berman, R., Boucksom, S., Guedj, V., Zeriahi, A.: A variational approach to complex Monge–Ampère equations. *Publ. Math. Inst. Hautes Études Sci.* **117**, 179–245 (2013)
7. Boucksom, S., Eyssidieux, P., Guedj, V., Zeriahi, A.: Monge–Ampère equations in big cohomology classes. *Acta Math.* **205**, 199–262 (2010)
8. Bedford, E., Taylor, B.A.: A new capacity for plurisubharmonic functions. *Acta Math.* **149**, 1–40 (1982)
9. Bedford, E., Taylor, B.A.: Fine topology, Shilov boundary and $(dd^c)^n$. *J. Funct. Anal.* **72**, 225–251 (1987)
10. Błocki, Z.: Uniqueness and stability for the complex Monge–Ampère equation on compact Kähler manifolds. *Indiana Univ. Math. J.* **52**, 1697–1701 (2003)
11. Błocki, Z., Kołodziej, S.: On regularization of plurisubharmonic functions on manifolds. *Proc. Am. Math. Soc.* **135**, 2089–2093 (2007)
12. Boucksom, S., Demailly, J.-P., Păun, M., Peternell, T.: The pseudo-effective cone of a compact Kähler manifold and varieties of negative Kodaira dimension. *J. Algebraic Geom.* **22**, 201–248 (2013)
13. Boucksom, S., Favre, C., Jonsson, M.: Solution to a non-Archimedean Monge–Ampère equation. *J. Am. Math. Soc.* **28**, 617–667 (2015)
14. Boucksom, S., Favre, C., Jonsson, M.: Singular semipositive metrics in non-Archimedean geometry. *J. Algebr. Geom.* **25**(1), 77–139 (2016)
15. Boucksom, S., Hisamoto, T., Jonsson, M.: Uniform K-stability, Duistermaat–Heckman measures and singularities of pairs. [arXiv:1504.06568](https://arxiv.org/abs/1504.06568)
16. Boucksom, S., Hisamoto, T., Jonsson, M.: Uniform K-stability and asymptotics of energy functionals in Kähler geometry. [arXiv:1603.01026](https://arxiv.org/abs/1603.01026)
17. Boucksom, S., Jonsson, M.: Singular semipositive metrics on line bundles on varieties over trivially valued fields. (in preparation)
18. Burgos Gil, J.I., Philippon, P., Sombra, M.: Arithmetic geometry of toric varieties. Metrics, measures, and heights. *Astérisque* **360**, vi+222 (2014)
19. Cegrell, U.: Pluricomplex energy. *Acta Math.* **180**, 187–217 (1998)
20. Chambert-Loir, A.: Mesures et équidistribution sur les espaces de Berkovich. *J. Reine Angew. Math.* **595**, 215–235 (2006)
21. Chambert-Loir, A.: Heights and measures on analytic spaces. A survey of recent results, and some remarks. In: *Motivic Integration and its Interactions with Model Theory and Non-Archimedean Geometry. Volume II*. London Mathematical Society Lecture Note Series, vol. 384. Cambridge University Press, Cambridge (2011)
22. Chambert-Loir, A., Ducros, A.: Formes différentielles réelles et courants sur les espaces de Berkovich. [arXiv:1204.6277](https://arxiv.org/abs/1204.6277)
23. Demailly, J.-P.: Regularization of closed positive currents and intersection theory. *J. Algebraic Geom.* **1**(3), 361–409 (1992)
24. Dinew, S.: Uniqueness in $\mathcal{E}(X, \omega)$. *J. Funct. Anal.* **256**, 2113–2122 (2009)
25. Eyssidieux, P., Guedj, V., Zeriahi, A.: Singular Kähler–Einstein metrics. *J. Am. Math. Soc.* **22**, 607–639 (2009)
26. Favre, C., Jonsson, M.: *The Valutive Tree*. Lecture Notes in Mathematics, vol. 1853. Springer, Berlin (2004)
27. Fulton, W.: *Introduction to Toric Varieties*. Annals of Mathematics Studies, vol. 131. Princeton University Press, Princeton, NJ (1993)
28. Gubler, W.: Local heights of subvarieties over non-Archimedean fields. *J. Reine Angew. Math.* **498**, 61–113 (1998)

29. Gubler, W.: Non-archimedean canonical measures on abelian varieties. *Compos. Math.* **146**, 683–730 (2010)
30. Gubler, W.: A guide to tropicalizations. In: *Algebraic and Combinatorial Aspects of Tropical Geometry*. Contemporary Mathematics, vol. 589, pp. 125–189. American Mathematical Society, Providence, RI (2013)
31. Gubler, W.: Forms and currents on the analytification of an algebraic variety (after Chambert-Loir and Ducros). [arXiv:1303.7364](https://arxiv.org/abs/1303.7364)
32. Gubler, W., Künnemann, K.: A tropical approach to non-Archimedean Arakelov theory. [arXiv:1406.7637](https://arxiv.org/abs/1406.7637)
33. Guedj, V., Zeriahi, A.: Intrinsic capacities on compact Kähler manifolds. *J. Geom. Anal.* **15**, 607–639 (2005)
34. Guedj, V., Zeriahi, A.: The weighted Monge-Ampère energy of quasiplurisubharmonic functions. *J. Funct. Anal.* **250**, 442–482 (2007)
35. Jonsson, M.: Degenerations of amoebae and Berkovich spaces. *Math. Ann.* **364**, 293–311 (2016)
36. Kempf, G., Knudsen, F., Mumford, D., Saint-Donat, B.: *Toroidal Embeddings. I*. Lecture Notes in Mathematics, vol. 339. Springer, Berlin (1973)
37. Kołodziej, S.: The complex Monge-Ampère equation. *Acta Math.* **180**, 69–117 (1998)
38. Kołodziej, S.: The Monge-Ampère equation on compact Kähler manifolds. *Indiana Univ. Math. J.* **52**, 667–686 (2003)
39. Kontsevich, M., Soibelman, Y.: Affine structures and non-Archimedean analytic spaces. In: *The Unity of Mathematics*. Progress in Mathematics, vol. 244, pp. 321–385. Birkhäuser, Boston (2006)
40. Kontsevich, M., Tschinkel, Y.: Non-Archimedean Kähler geometry. Unpublished (2000)
41. Lagerberg, A.: Super currents and tropical geometry. *Math. Z.* **270**, 1011–1050 (2012)
42. Lelong, P.: Les objets souples de l’analyse complexe. *Expo. Math.* **3**, 149–164 (1985)
43. Liu, Y.: A non-Archimedean analogue of Calabi-Yau theorem for totally degenerate abelian varieties. *J. Diff. Geom.* **89**, 87–110 (2011)
44. Mustață, M., Nicaise, J.: Weight functions on non-archimedean analytic spaces and the Kontsevich-Soibelman skeleton. *Algebr. Geom.* **2**, 365–404 (2015)
45. Nicaise, J., Xu, C.: The essential skeleton of a degeneration of algebraic varieties. *Am. J. Math.* [arXiv:1307.4041](https://arxiv.org/abs/1307.4041) (to appear)
46. Rauch, J., Taylor, B.A.: The Dirichlet problem for the multidimensional Monge-Ampère equation. *Rocky Mountain J. Math.* **7**(2), 345–364 (1977)
47. Thuillier, A.: Théorie du potentiel sur les courbes en géométrie analytique non archimédienne. Applications à la théorie d’Arakelov. Ph.D. Thesis (2005)
48. Yau, S.T.: On the Ricci curvature of a compact Kähler manifold and the complex Monge-Ampère equation. *Commun. Pure Appl. Math.* **31**, 339–411 (1978)
49. Yu, T.Y.: Gromov compactness in non-Archimedean analytic geometry. *J. Reine Angew. Math.* [arXiv:1401.6452](https://arxiv.org/abs/1401.6452) (to appear)
50. Yuan, X., Zhang, S.-W.: The arithmetic Hodge index theorem for adelic line bundles I. [arXiv:1304.3538](https://arxiv.org/abs/1304.3538)
51. Zhang, S.-W.: Small points and adelic metrics. *J. Algebraic Geom.* **4**, 281–300 (1995)

Convergence Polygons for Connections on Nonarchimedean Curves

Kiran S. Kedlaya

In classical analysis, one builds the catalog of special functions by repeatedly adjoining solutions of differential equations whose coefficients are previously known functions. Consequently, the properties of special functions depend crucially on the basic properties of ordinary differential equations. This naturally led to the study of formal differential equations, as in the seminal work of Turrittin [165]; this may be viewed retroactively as a theory of differential equations over a trivially valued field. After the introduction of p -adic analysis in the early twentieth century, there began to be corresponding interest in solutions of p -adic differential equations; however, aside from some isolated instances (e.g., the proof of the Nagell–Lutz theorem; see Theorem 3.4), a unified theory of p -adic ordinary differential equations did not emerge until the pioneering work of Dwork on the relationship between p -adic special functions and the zeta functions of algebraic varieties over finite fields (e.g., see [57, 58]). At that point, serious attention began to be devoted to a serious discrepancy between the p -adic and complex-analytic theories: on an open p -adic disc, a nonsingular differential equation can have a formal solution which does not converge in the entire disc (e.g., the exponential series). One is thus led to quantify the convergence of power series solutions of differential equations involving rational functions over a nonarchimedean field; this was originally done by Dwork in terms of the *generic radius of convergence* [59]. This and more refined invariants were studied by numerous authors during the half-century following Dwork’s initial work, as documented in the author’s book [92].

At around the time that [92] was published, a new perspective was introduced by Baldassarri [13] (and partly anticipated in prior unpublished work of Baldassarri and Di Vizio [15]) which makes full use of Berkovich’s theory of nonarchimedean analytic spaces. Given a differential equation as above, or more generally a

K.S. Kedlaya (✉)

Department of Mathematics, University of California, San Diego, La Jolla, CA 92093-0112, USA
e-mail: kedlaya@ucsd.edu

connection on a curve over a nonarchimedean field, one can define an invariant called the *convergence polygon*; this is a function from the underlying Berkovich topological space of the curve into a space of Newton polygons, which measures the convergence of formal horizontal sections and is well-behaved with respect to both the topology and the piecewise linear structure on the Berkovich space. One can translate much of the prior theory of p -adic differential equations into (deceptively) simple statements about the behavior of the convergence polygon; this process was carried out in a series of papers by Poineau and Pulita [132, 134, 135, 140], as supplemented by work of this author [97] and upcoming joint work with Baldassarri [16].

In this paper, we present the basic theorems on the convergence polygon, which provide a number of combinatorial constraints that may be used to extract information about convergence of formal horizontal sections at one point from corresponding information at other points. We include numerous examples to illustrate some typical behaviors of the convergence polygon. We also indicate some relationships between convergence polygons, index formulas for de Rham cohomology, and the geometry of finite morphisms, paying special attention to the case of cyclic p -power coverings with p equal to the residual characteristic. This case is closely linked with the Oort lifting problem for Galois covers of curves in characteristic p , and some combinatorial constructions arising in that problem turn out to be closely related to convergence polygons. There are additional applications to the study of integrable connections on higher-dimensional nonarchimedean analytic spaces, both in the cases of zero residual characteristic [94] and positive residual characteristic [95], but we do not pursue these applications here.

To streamline the exposition, we make no attempt to indicate the techniques of proof underlying our main results; in some cases, quite sophisticated arguments are required. We limit ourselves to saying that the basic tools are developed in a self-contained fashion in [92] (but without reference to Berkovich spaces), and the other aforementioned results are obtained by combining results from [92] in an intricate manner. (Some exceptions are made for results which do not occur in any existing paper; their proofs are relegated to an Appendix.) We also restrict generality by considering only proper curves, even though many of the results we discuss can be formulated for open curves, possibly of infinite genus.

At the suggestion of the referee, we include (see Appendix 2) a *thematic bibliography* in the style of [147] listing additional references germane to the topics discussed in the article. We also include references for the following topics which we have omitted out of space considerations, even though in practice they are thoroughly entangled with the proofs of the results which we do mention.

- *Decomposition theorems*, which give splittings of connections analogous to the factorizations of polynomials given by various forms of Hensel's lemma.
- *Monodromy theorems*, which provide structure theorems for connections at the expense of passing from a given space to a finite cover.
- *Logarithmic growth*, i.e., secondary terms in the measurement of convergence of horizontal sections.

1 Newton Polygons

As setup for our definition of convergence polygons, we fix some conventions regarding Newton polygons.

Definition 1.1. For n a positive integer, let $\mathcal{P}[0, n]$ be the set of continuous functions $\mathcal{N} : [0, n] \rightarrow \mathbb{R}$ satisfying the following conditions.

- (a) We have $\mathcal{N}(0) = 0$.
- (b) For $i = 1, \dots, n$, the restriction of \mathcal{N} to $[i - 1, i]$ is affine.

For $i = 1, \dots, n$, we write $h_i : \mathcal{P}[0, n] \rightarrow \mathbb{R}$ for the function $\mathcal{N} \mapsto \mathcal{N}(i)$; we call $h_i(\mathcal{N})$ the i -th *height* of \mathcal{N} . The product map $h_1 \times \dots \times h_n : \mathcal{P}[0, n] \rightarrow \mathbb{R}^n$ is a bijection, using which we equip $\mathcal{P}[0, n]$ with a topology and an integral piecewise linear structure. We sometimes refer to h_n simply as h and call it the *total height*.

Definition 1.2. Let $\mathcal{NP}[0, n]$ be the subset of $\mathcal{P}[0, n]$ consisting of *concave* functions. For $i = 1, \dots, n$, we write $s_i : \mathcal{P}[0, n] \rightarrow \mathbb{R}$ for the function $\mathcal{N} \mapsto \mathcal{N}(i) - \mathcal{N}(i - 1)$; we call $s_i(\mathcal{N})$ the i -th *slope* of \mathcal{N} . For $\mathcal{N} \in \mathcal{NP}[0, n]$, we have $s_1(\mathcal{N}) \geq \dots \geq s_n(\mathcal{N})$.

Definition 1.3. Let $I \subseteq \mathbb{R}$ be a closed interval. A function $\mathcal{N} : I \rightarrow \mathcal{NP}[0, n]$ is *affine* if it has the form $\mathcal{N}(t) = \mathcal{N}_0 + t\mathcal{N}_1$ for some $\mathcal{N}_0, \mathcal{N}_1 \in \mathcal{P}[0, n]$. In this case, we call \mathcal{N}_1 the *slope* of \mathcal{N} . We say that \mathcal{N} has *integral derivative* if $\mathcal{N}_1(i) \in \mathbb{Z}$ for $i = n$ and for each $i \in \{1, \dots, n - 1\}$ such that for all t in the interior of I , $\mathcal{N}(t)$ has a change of slope at i . This implies that the graph of \mathcal{N}_1 has vertices only at lattice points, but not conversely. (It would be natural to use the terminology *integral slope*, but we avoid this terminology to alleviate confusion with Definition 1.2.)

2 PL Structures on Berkovich Curves

We next recall the canonical piecewise linear structure on a Berkovich curve (e.g., see [9]).

Hypothesis 2.1. For the rest of this paper, let K be a *nonarchimedean field*, i.e., a field complete with respect to a nonarchimedean absolute value; let X be a smooth, proper, geometrically connected curve over K ; let Z be a finite set of closed points of X ; and put $U = X - Z$.

Convention 2.2. Whenever we view \mathbb{Q}_p as a nonarchimedean field, we normalize the p -adic absolute value so that $|p| = p^{-1}$.

Remark 2.3. Recall that the points of the Berkovich analytification X^{an} may be identified with equivalence classes of pairs (L, x) in which L is a nonarchimedean field over K and x is an element of $X(L)$, where the equivalence relation is generated by relations of the form $(L, x) \sim (L', x')$ where x' is the restriction of x along a

continuous K -algebra homomorphism $L \rightarrow L'$. As is customary, we classify points of X^{an} into types 1, 2, 3, 4 (e.g., see [97, Proposition 4.2.7]). To lighten notation, we identify Z with Z^{an} , which is a finite subset of X^{an} consisting of type 1 points.

Definition 2.4. For $\rho > 0$, let x_ρ denote the generic point of the disc $|z| \leq \rho$ in \mathbb{P}_K^1 . A *segment* in X^{an} is a closed subspace S homeomorphic to a closed interval for which there exist an open subspace V of X^{an} , a choice of values $0 \leq \alpha < \beta \leq +\infty$, and an isomorphism of V with $\{z \in \mathbb{P}_K^{1,\text{an}} : \alpha < |z| < \beta\}$ identifying the interior of S with $\{x_\rho : \rho \in (\alpha, \beta)\}$. A *virtual segment* in X^{an} is a connected closed subspace whose base extension to some finite extension of K is a disjoint union of segments.

A *strict skeleton* in X^{an} is a subspace Γ containing Z^{an} equipped with a homeomorphism to a finite connected graph, such that each vertex of the graph corresponds to either a point of Z or a point of type 2, and each edge corresponds to a virtual segment, and X^{an} retracts continuously onto Γ . Using either tropicalizations or semistable models, one may realize X^{an} as the inverse limit of its strict skeletons; again, see [9] for a detailed discussion.

Definition 2.5. Note that

$$\chi(U) = 2 - 2g(X) - \text{length}(Z),$$

so $\chi(U) \leq 0$ if and only if either $g(X) \geq 1$ or $\text{length}(Z) \geq 2$. In this case, there is a unique minimal strict skeleton in X^{an} , which we denote $\Gamma_{X,Z}$. Explicitly, if K is algebraically closed, then the underlying set of $\Gamma_{X,Z}$ is the complement in X^{an} of the union of all open discs in U^{an} ; for general K , the underlying set of $\Gamma_{X,Z}$ is the image under restriction of the minimal strict skeleton in X_L^{an} for L a completed algebraic closure of K . In particular, if K' is the completion of an algebraic extension of K , then the minimal strict skeleton in $X_{K'}^{\text{an}}$ is the inverse image of $\Gamma_{X,Z}$ in $X_{K'}^{\text{an}}$. However, the corresponding statement for a general nonarchimedean field extension K' of K is false; see Definition 8.5 for a related phenomenon.

3 Convergence Polygons: Projective Line

We next introduce the concept of the convergence polygon associated with a differential equation on \mathbb{P}^1 .

Hypothesis 3.1. For the rest of this paper, we assume that the nonarchimedean field K is of characteristic 0, as otherwise the study of differential operators on K -algebras has a markedly different flavor (for instance, any derivation on a ring R of characteristic $p > 0$ has the subring of p -th powers in its kernel). By contrast, the residue characteristic of K , which we call p , may be either 0 or positive unless otherwise specified (e.g., if we refer to \mathbb{Q}_p , then we implicitly require $p > 0$).

Hypothesis 3.2. For the rest of Sect. 3, take $X = \mathbb{P}_K^1$, assume $\infty \in Z$, and consider the differential equation

$$y^{(n)} + f_{n-1}(z)y^{(n-1)} + \dots + f_0(z)y = 0 \tag{1}$$

for some rational functions $f_0, \dots, f_{n-1} \in K(z)$ with poles only within Z . If $Z = \{\infty\}$, let m be the dimension of the K -vector space of entire solutions of (1); otherwise, take $m = 0$.

Definition 3.3. For any nonarchimedean field L over K and any $x \in U(L)$, let S_x be the set of formal solutions of (1) with $y \in L[[z - x]]$. By interpreting (1) as a linear recurrence relation of order n on the coefficients of a power series, we see that every list of n initial conditions at $z = x$ corresponds to a unique formal solution; that is, the composition

$$S_x \rightarrow L[[z - x]] \rightarrow L[[z - x]]/(z - x)^n$$

is a bijection. In particular, S_x is an L -vector space of dimension n .

Theorem 3.4 (*p*-Adic Cauchy Theorem). *Each element of S_x has a positive radius of convergence.*

Proof. This result was originally proved by Lutz [106, Théorème IV] somewhat before the emergence of the general theory of p -adic differential equations; Lutz used it as a lemma in her proof of the Nagell–Lutz theorem on the integrality of torsion points on rational elliptic curves. One can give several independent proofs using the modern theory; see [92, Propositions 9.3.3, 18.1.1]. \square

Definition 3.5. For $i = 1, \dots, n - m$, choose $s_i(x) \in \mathbb{R}$ so that $e^{-s_i(x)}$ is the supremum of the set of $\rho > 0$ such that U^{an} contains the open disc $|z - x| < \rho$ and S_x contains $n - i + 1$ linearly independent elements convergent on this disc. Note that this set is nonempty by Theorem 3.4 and bounded above by the definition of m , so the definition makes sense. In particular, $s_1(x)$ is the joint radius of convergence of all of the elements of S_x , while $s_{n-m}(x)$ is the maximum finite radius of convergence of a nonzero element of S_x .

Since $s_1(x) \geq \dots \geq s_{n-m}(x)$, the $s_i(x)$ are the slopes of a polygon $\mathcal{N}_z(x) \in \mathcal{NP}[0, n - m]$, which we call the *convergence polygon* of (1) at x . (We include z in the notation to remind ourselves that \mathcal{N}_z depends on the choice of the coordinate z of X .) This construction is compatible with base change: if L' is a nonarchimedean field containing L and x' is the image of x in $U(L')$, then $\mathcal{N}_z(x) = \mathcal{N}_z(x')$. Consequently, we obtain a well-defined function $\mathcal{N}_z : U^{\text{an}} \rightarrow \mathcal{NP}[0, n - m]$.

Definition 3.6. Suppose that $Z \neq \{\infty\}$. By definition, $e^{-s_1(\mathcal{N}_z(x))}$ can never exceed the largest value of ρ for which the disc $|z - x| < \rho$ does not meet Z . When equality occurs, we say that (1) satisfies the *Robba condition* at x .

Theorem 3.7. *The function $\mathcal{N}_z : U^{\text{an}} \rightarrow \mathcal{NP}[0, n - m]$ is continuous; more precisely, it factors through the retraction of $\mathbb{P}_K^{1, \text{an}}$ onto some strict skeleton Γ , and the restriction of \mathcal{N}_z to each edge of Γ is affine with integral derivative.*

Proof. See [140] or [97] or [16]. □

One can say quite a bit more, but for this it is easier to shift to a coordinate-free interpretation, which also works for more general curves; see Sect. 5.

4 A Gallery of Examples

To help the reader develop some intuition, we collect a few illustrative examples of convergence polygons. Throughout Sect. 4, retain Hypothesis 3.1.

Example 4.1. Take $K = \mathbb{Q}_p$, $Z = \{\infty\}$, and consider the differential equation $y' - y = 0$. The formal solutions of this equation with $y \in L[[z - x]]$ are the scalar multiples of the exponential series

$$\exp(z - x) = \sum_{i=0}^{\infty} \frac{(z - x)^i}{i!},$$

which has radius of convergence $p^{-1/(p-1)}$. Consequently,

$$s_1(\mathcal{N}_z(x)) = \frac{1}{p-1} \log p;$$

in particular, \mathcal{N}_z is constant on U^{an} .

In this next example, we illustrate the effect of changing Z on the convergence polygon.

Example 4.2. Set notation as in Example 4.1, except now with $Z = \{0, \infty\}$. In this case we have

$$s_1(\mathcal{N}_z(x)) = \max \left\{ -\log |x|, \frac{1}{(p-1)} \log p \right\}.$$

In particular, \mathcal{N}_z factors through the retraction of $\mathbb{P}_K^{1, \text{an}}$ onto the path from 0 to ∞ . For $x \in U^{\text{an}}$, the Robba condition holds at x if and only if $|x| \leq p^{-1/(p-1)}$.

Example 4.3. Take $K = \mathbb{Q}_p$, $Z = \{0, \infty\}$, and consider the differential equation $y' - \frac{1}{p}z^{-1}y = 0$. The formal solutions of this equation with $y \in L[[z - x]]$ are the scalar multiples of the binomial series

$$\sum_{i=0}^{\infty} \binom{1/p}{i} x^{-i+p^{-1}} (z - x)^i,$$

which has radius of convergence $p^{-p/(p-1)}|x|$. Consequently,

$$s_1(\mathcal{N}_z(x)) = \frac{p}{p-1} \log p - \log |x|,$$

so again \mathcal{N}_z factors through the retraction of $\mathbb{P}_K^{1,\text{an}}$ onto the path from 0 to ∞ . In this case, the Robba condition holds nowhere.

Example 4.4. Assume $p > 2$, take $K = \mathbb{Q}_p$, $Z = \{0, \infty\}$, and consider the Bessel differential equation (with parameter 0)

$$y'' + z^{-1}y' + y = 0.$$

This example was studied by Dwork [61], who showed that

$$s_1(\mathcal{N}_z(x)) = s_2(\mathcal{N}_z(x)) = \max \left\{ -\log |x|, \frac{1}{p-1} \log p \right\}.$$

Again, \mathcal{N}_z factors through the retraction of $\mathbb{P}_K^{1,\text{an}}$ onto the path from 0 to ∞ . As in Example 4.2, for $x \in U^{\text{an}}$, the Robba condition holds at x if and only if $|x| \leq p^{-1/(p-1)}$.

Our next example illustrates a typical effect of varying a parameter.

Example 4.5. Let K be an extension of \mathbb{Q}_p , take $Z = \{0, \infty\}$, and consider the differential equation $y' - \lambda z^{-1}y = 0$ for some $\lambda \in K$ (the case $\lambda = 1/p$ being Example 4.3). Then

$$s_1(\mathcal{N}_z(x)) = c + \frac{1}{p-1} \log p - \log |x|$$

where c is a continuous function of

$$c_0 = \min\{|\lambda - t| : t \in \mathbb{Z}_p\};$$

namely, by [66, Proposition IV.7.3] we have

$$c = \begin{cases} \log c_0 & \text{if } c_0 \geq 1 \\ -\frac{p^m-1}{(p-1)p^m} \log p + \frac{1}{p^{m+1}} \log(p^m c_0) & \text{if } p^{-m-1} \leq c_0 \leq p^{-m} \quad (m = 0, 1, \dots) \\ -\frac{1}{p-1} \log p & \text{if } c_0 = 0. \end{cases}$$

In particular, the Robba condition holds everywhere if $\lambda \in \mathbb{Z}_p$ and nowhere otherwise. In either case, \mathcal{N}_z factors through the retraction of $\mathbb{P}_K^{1,\text{an}}$ onto the path from 0 to ∞ .

Example 4.6. Take $K = \mathbb{Q}_p$, $Z = \{\infty\}$, and consider the differential equation $y' - az^{a-1}y = 0$ for some positive integer a not divisible by p (the case $a = 1$ being Example 4.1). The formal solutions of this equation are the scalar multiples of

$$\exp(z^a - x^a).$$

This series converges in the region where $|z^a - x^a| < p^{-1/(p-1)}$; consequently,

$$s_1(\mathcal{N}_z(x)) = \max \left\{ \frac{1}{p-1} \log p + (a-1) \log |x|, \frac{1}{a(p-1)} \log p \right\}.$$

In this case, \mathcal{N}_z factors through the retraction onto the path from $x_{p^{-1/(p-1)}}$ to ∞ .

Example 4.7. Take $K = \mathbb{C}((t))$ (so $p = 0$), $Z = \{\infty\}$, and consider the differential equation

$$y''' + zy'' + y = 0.$$

It can be shown that

$$s_1(\mathcal{N}_z(x)) = \max \{0, \log |x|\}, s_2(\mathcal{N}_z(x)) = s_3(\mathcal{N}_z(x)) = \min \left\{ 0, -\frac{1}{2} \log |x| \right\}.$$

In this case, \mathcal{N}_z factors through the retraction onto the path from x_1 to ∞ . Note that this provides an example where the slopes of $\mathcal{N}_z(x)$ are not bounded below uniformly on $(\mathbb{P}_K^1 - Z)^{\text{an}}$; that is, as x approaches ∞ , two linearly independent local solutions have radii of convergence growing without bound, but these local solutions do not patch together.

Example 4.8. Assume $p > 2$, take $K = \mathbb{Q}_p$, $Z = \{0, 1, \infty\}$, and consider the Gaussian hypergeometric differential equation

$$y'' + \frac{(1-2z)}{z(1-z)}y' - \frac{1}{4z(1-z)}y = 0.$$

This example was originally studied by Dwork [58] due to its relationship with the zeta functions of elliptic curves. Using Dwork's calculations, it can be shown that

$$s_1(\mathcal{N}_z(x)) = s_2(\mathcal{N}_z(x)) = \max \{ \log |x|, -\log |x|, -\log |x-1| \}.$$

In this case, \mathcal{N}_z factors through the retraction from $\mathbb{P}_K^{1,\text{an}}$ onto the union of the paths from 0 to ∞ and from 1 to ∞ , and the Robba condition holds everywhere.

Remark 4.9. One can compute additional examples of convergence polygons associated with first-order differential equations using an explicit formula for the radius of convergence at a point, due to Christol–Pulita. This result was originally reported in [36] but with an error in the formula; for a corrected statement, see [141, Introduction, Théorème 5].

5 Convergence Polygons: General Curves

We now describe an analogue of the convergence polygon in a more geometric setting.

Hypothesis 5.1. Throughout Sect. 5, assume that $\chi(U) \leq 0$, i.e., either $g(X) \geq 1$ or $\text{length}(Z) \geq 2$. Let \mathcal{E} be a vector bundle on U of rank n equipped with a connection $\nabla : \mathcal{E} \rightarrow \mathcal{E} \otimes_{\mathcal{O}_U} \Omega_{U/K}$. As is typical, we describe sections of \mathcal{E} in the kernel of ∇ as being *horizontal*.

Remark 5.2. For the results in this section, one could also allow X to be an analytic curve which is compact but not necessarily proper. To simplify the discussion, we omit this level of generality; the general results can be found in any of [16, 97, 132].

Definition 5.3. Let L be a nonarchimedean field containing K and choose $x \in U(L)$, which we identify canonically with a point of U_L^{an} . Since X is smooth, U_L^{an} contains a neighborhood of x isomorphic to an open disc over L . Thanks to our restrictions on X and Z , the union U_x of all such neighborhoods in U_L^{an} is itself isomorphic to an open disc over L . The construction is compatible with base change in the following sense: if L' is a nonarchimedean field containing L , then $U_{x,L'}$ is the union of all neighborhoods of x in $U_{L'}^{\text{an}}$ isomorphic to an open disc over L' . For each $\rho \in (0, 1]$, let $U_{x,\rho}$ be the open disc of radius ρ centered at x within U_x (normalized so that $U_{x,1} = U_x$).

Let $\widehat{\mathcal{O}}_{X_L,x}$ denote the completed local ring of X_L at x ; it is abstractly a power series ring in one variable over L . Let \mathcal{E}_x denote the pullback of \mathcal{E} to $\widehat{\mathcal{O}}_{X_L,x}$, equipped with the induced connection. One checks easily that \mathcal{E}_x is a trivial differential module; more precisely, the space $\ker(\nabla, \mathcal{E}_x)$ is an n -dimensional vector space over L and the natural map

$$\ker(\nabla, \mathcal{E}_x) \otimes_L \widehat{\mathcal{O}}_{X_L,x} \rightarrow \mathcal{E}_x$$

is an isomorphism.

For $i = 1, \dots, n$, choose $s_i(x) \in [0, +\infty)$ so that $e^{-s_i(x)}$ is the supremum of the set of $\rho \in (0, 1]$ such that \mathcal{E}_x contains $n - i + 1$ linearly independent sections convergent on $U_{x,\rho}$. Again, this set of such ρ is nonempty by Theorem 3.4. Since $s_1(x) \geq \dots \geq s_n(x)$, the $s_i(x)$ are the slopes of a polygon $\mathcal{N}(x) \in \mathcal{NP}[0, n]$, which we call the *convergence polygon* of \mathcal{E} at x . Again, the construction is compatible with base change, so it induces a well-defined function $\mathcal{N} : U^{\text{an}} \rightarrow \mathcal{NP}[0, n]$.

Definition 5.4. For $x \in U^{\text{an}}$, we say that \mathcal{E} satisfies the *Robba condition* at x if $\mathcal{N}(x)$ is the zero polygon.

We have the following analogue of Theorem 3.7.

Theorem 5.5. *The function $\mathcal{N} : (X - Z)^{\text{an}} \rightarrow \mathcal{NP}[0, n]$ is continuous. More precisely, there exists a strict skeleton Γ such that \mathcal{N} factors through the retraction of X^{an} onto Γ , and the restriction of \mathcal{N} to each edge of Γ is affine with integral derivative.*

Proof. See [132] or [97] or [16] (and Remark 5.6). \square

Remark 5.6. It is slightly inaccurate to attribute Theorem 5.5 to [132] or [16], as the results proved therein are slightly weaker: they require an uncontrolled base extension on K , which creates more options for the strict skeleton Γ . In particular, Theorem 5.5 as stated implies that \mathcal{N} is locally constant around any point of type 4, which cannot be established using the methods of [132] or [16]; one instead requires some dedicated arguments found only in [97]. These extra arguments are crucial for the applications of Theorem 5.5 in the contexts described in [94, 95].

Remark 5.7. Suppose that $X = \mathbb{P}_K^1$ and $\infty \in Z$. Given a differential equation as in (1), we can construct an associated connection \mathcal{E} of rank n whose underlying vector bundle is free on the basis $\mathbf{e}_1, \dots, \mathbf{e}_n$ and whose action of ∇ is given by

$$\begin{aligned} \nabla(\mathbf{e}_1) &= f_0(z)\mathbf{e}_n \\ \nabla(\mathbf{e}_2) &= f_1(z)\mathbf{e}_n - \mathbf{e}_1 \\ &\vdots \\ \nabla(\mathbf{e}_{n-1}) &= f_{n-2}(z)\mathbf{e}_n - \mathbf{e}_{n-2} \\ \nabla(\mathbf{e}_n) &= f_{n-1}(z)\mathbf{e}_n - \mathbf{e}_{n-1}. \end{aligned}$$

A section of \mathcal{E} is then horizontal if and only if it has the form $y\mathbf{e}_1 + y'\mathbf{e}_2 + \dots + y^{(n-1)}\mathbf{e}_n$ where y is a solution of (1). If $\text{length}(Z) \geq 2$, each of \mathcal{N}_z and \mathcal{N} can be computed in terms of the other; this amounts to changing the normalization of certain discs. In particular, the statements of Theorems 3.7 and 5.5 in this case are equivalent.

If $\text{length}(Z) = 1$, we cannot define \mathcal{N} as above. However, if K is nontrivially valued, one can recover the properties of \mathcal{N}_z by considering \mathcal{N} with Z replaced by $Z \cup \{x\}$ for some $x \in U(K)$ with $|x|$ sufficiently large (namely, larger than the radius of convergence of any nonentire formal solution at 0). We refer to [16] for further details.

Remark 5.8. One can extend Remark 5.7 by defining \mathcal{N}_z in the case where $X = \mathbb{P}_K^1$ and $\infty \in Z$, and using Theorem 5.5 to establish an analogue of Theorem 3.7. With this modification, we still do not define either \mathcal{N} or \mathcal{N}_z in the case where $X = \mathbb{P}_K^1$ and $Z = \emptyset$, but this case is completely trivial: the vector bundle \mathcal{E} must admit a basis of horizontal sections (see [16]).

Theorem 5.9. *Suppose that $x \in \Gamma \cap U^{\text{an}}$ is the generic point of an open disc D contained in X and the Robba condition holds at x .*

- (a) *If $D \cap Z = \emptyset$, then the restriction of \mathcal{E} to D is trivial (i.e., it admits a basis of horizontal sections).*
- (b) *If $D \cap Z$ consists of a single point z at which ∇ is regular, then the Robba condition holds on $D - \{z\}$.*

Proof. Part (a) is a special case of the Dwork transfer theorem; see, for instance, [92, Theorem 9.6.1]. Part (b) follows as in the proof of [92, Theorem 13.7.1]. \square

Remark 5.10. Theorem 5.9(b) is a variant of a result of Christol, which has a slightly stronger hypothesis and a slightly stronger conclusion. In Christol’s result, one must assume either that $p = 0$ or that $p > 0$, and the pairwise differences between the exponents of ∇ at z are not p -adic Liouville numbers (see Example 7.19). One however gets the stronger conclusion that the “formal solution matrix” of ∇ at z converges on all of D . Both parts of Theorem 5.9 are examples of *transfer theorems*, which can be viewed as transferring convergence information from one disc to another.

Remark 5.11. Let k be the residue field of K . Suppose that $X = \mathbb{P}_K^1$, ∇ is regular everywhere, and the reduction map from Z to \mathbb{P}_k^1 is injective. Then Theorem 5.9 implies that if the Robba condition holds at x_1 , then it holds on all of U^{an} . For instance, this is the case for (the connection associated via Remark 5.7 to) the hypergeometric equation considered in Example 4.8; more generally, it holds for the hypergeometric equation

$$y'' - \frac{c - (a + b + 1)z}{z(1 - z)}y' - \frac{ab}{z(1 - z)}y = 0$$

if and only if $a, b, c \in \mathbb{Z}_p$ (the case of Example 4.8 being $a = b = 1/2, c = 1$). This example and Example 4.5, taken together, suggest that for a general differential equation with one or more accessory parameters, the Robba condition at a fixed point is likely to be of a “fractal” nature in these parameters. For some additional examples with four singular points, see the work of Beukers [21].

Remark 5.12. One can also consider some modified versions of the convergence polygon. For instance, one might take $e^{-s_i(x)}$ to be the supremum of those $\rho \in (0, 1]$ such that the restriction of \mathcal{E} to $U_{x,\rho}$ splits off a trivial submodule of rank at least $n - i + 1$; the resulting convergence polygons will again satisfy Theorem 5.5. It may be that some modification of this kind can be used to eliminate some hypotheses on p -adic exponents, as in Theorem 7.23.

6 Derivatives of Convergence Polygons

We now take a closer look at the local variation of convergence polygons. Throughout Sect. 6, continue to retain Hypothesis 5.1.

Definition 6.1. For $x \in X^{\text{an}}$, a *branch* of X at x is a local connected component of $X - \{x\}$, that is, an element of the direct limit of $\pi_0(U - \{x\})$ as U runs over all neighborhoods of x in X . Depending on the type of x , the branches of X can be described as follows.

Type 1: A single branch.

Type 2: One branch corresponding to each closed point on the curve C_x (defined over the residue field of K) whose function field is the residue field of $\mathcal{H}(x)$.

Type 3: Two branches.

Type 4: One branch.

For each branch \vec{t} of X at x , by Theorem 5.5 we may define the *derivative* of \mathcal{N} along \vec{t} (away from x), as an element of $\mathcal{P}[0, n]$ with integral vertices; we denote this element by $\partial_{\vec{t}}(\mathcal{N})$. For x of type 1, we also denote this element by $\partial_x(\mathcal{N})$ since there is no ambiguity about the choice of the branch. We may similarly define $\partial_{\vec{t}}(h_i(\mathcal{N})) \in \mathbb{R}$ for $i = 1, \dots, n$, optionally omitting i in the case $i = n$; note that $\partial_{\vec{t}}(h(\mathcal{N})) \in \mathbb{Z}$.

Theorem 6.2. *For $z \in Z$, $-\partial_z(\mathcal{N})$ is the polygon associated with the Turrittin–Levelt–Hukuhara decomposition of \mathcal{E}_z (see [66, Chap. 4], [82, Sect. 11], or [92, Chap. 7]). In particular, this polygon belongs to $\mathcal{NP}[0, n]$, its slopes are all nonnegative, and its height equals the irregularity $\text{Irr}_z(\nabla)$ of ∇ at z .*

Proof. See [133, Sect. 5.7] and [134, Sect. 3.6], or see [16]. □

Corollary 6.3. *For $z \in Z$, \mathcal{N} extends continuously to a neighborhood of z if and only if ∇ has a regular singularity at z (i.e., its irregularity at z equals 0). In particular, \mathcal{N} extends continuously to all of X^{an} if and only if ∇ is everywhere regular.*

Remark 6.4. Using a similar technique, one can compute the asymptotic behavior of \mathcal{N} in a neighborhood of $z \in Z$ in terms of the “eigenvalues” occurring in the Turrittin–Levelt–Hukuhara decomposition of \mathcal{E}_z . For example, ∇ satisfies the Robba condition on some neighborhood of z if and only if ∇ is regular at z with all exponents in \mathbb{Z}_p .

Remark 6.5. In the complex-analytic setting, the decomposition of \mathcal{E}_z does not typically extend to any nonformal neighborhood of z ; this is related to the discrepancy between local indices in the algebraic and analytic categories (see Remark 7.8). Nonetheless, there are still some close links between the formal decomposition and the asymptotic behavior of local solutions on sectors at z (related to the theory of Stokes phenomena).

In the nonarchimedean setting, by contrast, the decomposition of \mathcal{E}_z always lifts to some nonformal neighborhood of z ; this follows from a theorem of Clark [46], which asserts (modulo some hypotheses as in Remark 7.18) that formal solutions of a connection at a (possibly irregular) singular point always converge in some punctured disc.

Theorem 6.6. *Assume $p = 0$. For $x \in U^{\text{an}}$ and \vec{t} a branch of X at x not pointing along $\Gamma_{X,Z}$, we have $\partial_{\vec{t}}(\mathcal{N}) \leq 0$.*

Proof. This requires somewhat technical arguments not present in the existing literature; see Theorem A.6. □

Remark 6.7. In the setting of Theorem 6.6, the statement that $\partial_{\vec{t}}(h_1(\mathcal{N})) \leq 0$ is equivalent to the Dwork transfer theorem (again see [92, Theorem 9.6.1]). In general, Theorem 6.6 is deduced by relating $\partial_{\vec{t}}(\mathcal{N})$ to local indices, as discussed in Sect. 7. This argument cannot work for $p > 0$ due to certain pathologies related to p -adic Liouville numbers (see Remark 7.18). We are hopeful that one can use a perturbation argument to deduce the analogue of Theorem 6.6 for $p > 0$, by reducing to the case where $\mathcal{N}(x)$ has no slopes equal to 0 and applying results of [92] (especially [92, Theorem 11.3.4]); however, a proof along these lines was not ready at the time of this writing.

Remark 6.8. By Theorem 5.5, for each $x \in U^{\text{an}}$, there exist only finitely many branches \vec{t} at x along which \mathcal{N} has nonzero slope. If x is of type 1 or 4, there are in fact no such branches. If x is of type 3, then the slopes along the two branches at x add up to 0.

7 Subharmonicity and Index

Using the piecewise affine structure of the convergence polygon, we formulate some additional properties, including local and global index formulas for de Rham cohomology. The local index formula is due to Poineau and Pulita [135], generalizing some partial results due to Robba [142, 143, 145, 146] and Christol–Mebkhout [40–42, 44]. Unfortunately, in the case $p > 0$ one is forced to interact with a fundamental pathology in the theory of p -adic differential equations, the effect of p -adic Liouville numbers; consequently, the global formula we derive here cannot be directly deduced from the local formula (see Remarks 7.14 and 7.18).

Hypothesis 7.1. Throughout Sect. 7, continue to retain Hypothesis 5.1, but assume in addition that K is algebraically closed. (Without this assumption, one can still formulate the results at the expense of having to keep track of some additional multiplicity factors.)

Definition 7.2. For $x \in U^{\text{an}}$, let $(\Delta\mathcal{N})_x \in \mathcal{P}[0, n]$ denote the sum of $\partial_{\vec{t}}(\mathcal{N})$ over all branches \vec{t} of X at x (oriented away from x); by Remark 6.8, this sum can only be nonzero when x is of type 2. Define the *Laplacian* of \mathcal{N} as the $\mathcal{P}[0, n]$ -valued measure $\Delta\mathcal{N}$ taking the continuous function $f : U^{\text{an}} \rightarrow \mathbb{R}$ to $\sum_{x \in U^{\text{an}}} f(x) (\Delta\mathcal{N})_x$. For $i = 1, \dots, n$, we may similarly define the multiplicities $(\Delta h_i(\mathcal{N}))_x \in \mathbb{R}$ and the Laplacian $\Delta h_i(\mathcal{N})$; we again omit the index i when it equals n .

Remark 7.3. The definition of the Laplacian can also be interpreted in the context of Thuillier’s potential theory [159], which applies more generally to functions which need not be piecewise affine.

Lemma 7.4. *We have*

$$\int \Delta h(\mathcal{N}) = \sum_{z \in Z} \text{Irr}_z(\nabla).$$

Proof. For e an edge of Γ , we may compute the slopes of \mathcal{N} along the two branches pointing into e from the endpoints of e ; these two slopes add up to 0. If we add up these slopes over all e , then regroup this sum by vertices, then the sum at each vertex $z \in Z$ equals $-\text{Irr}_z(\nabla)$ by Theorem 6.2, while the sum at each vertex $x \in U^{\text{an}}$ is the multiplicity of x in $\Delta \mathcal{N}$. This proves the claim. \square

Definition 7.5. For any open subset V of X^{an} , consider the complex

$$0 \rightarrow \mathcal{E} \xrightarrow{\nabla} \mathcal{E} \otimes \Omega \rightarrow 0$$

of sheaves, keeping in mind that if $V \cap Z \neq \emptyset$, then the sections over V are allowed to be meromorphic at $V \cap Z$ (but not to have essential singularities; see Remark 7.8). We define $\chi_{\text{dR}}(V, \mathcal{E})$ to be the index of the hypercohomology of this complex, i.e., the alternating sum of K -dimensions of the hypercohomology groups.

Lemma 7.6. *We have*

$$\chi_{\text{dR}}(X^{\text{an}}, \mathcal{E}) = n\chi(U) - \sum_{z \in Z} \text{Irr}_z(\nabla) = n(2 - 2g(X) - \text{length}(Z)) - \sum_{z \in Z} \text{Irr}_z(\nabla). \quad (2)$$

Proof. Let K_0 be a subfield of K which is finitely generated over \mathbb{Q} to which $X, Z, \mathcal{E}, \nabla$ can be descended. Then choose an embedding $K_0 \subset \mathbb{C}$ and let $X_{\mathbb{C}}$ be the base extension of the descent of X , again equipped with a meromorphic vector bundle \mathcal{E} and connection ∇ . Note that $\chi_{\text{dR}}(X^{\text{an}}, \mathcal{E})$ is computed by a spectral sequence in which one first computes the coherent cohomology of \mathcal{E} and $\mathcal{E} \otimes \Omega$ separately. By the GAGA principle both over \mathbb{C} [75, Exposé XII] and K [48, Example 3.2.6], these coherent cohomology groups can be computed equally well over any of X^{an} , X (or its descent to K_0), $X_{\mathbb{C}}$, or $X_{\mathbb{C}}^{\text{an}}$. Consequently, despite the fact that the connection is only K -linear rather than \mathcal{O} -linear, we may nonetheless conclude that $\chi_{\text{dR}}(X^{\text{an}}, \mathcal{E}) = \chi_{\text{dR}}(X_{\mathbb{C}}^{\text{an}}, \mathcal{E})$. (As an aside, this argument recovers a comparison theorem of Baldassarri [11].)

To compute $\chi_{\text{dR}}(X_{\mathbb{C}}^{\text{an}}, \mathcal{E})$, we may either appeal directly to [53, (6.21.1)] or argue directly as follows. Form a finite open covering $\{V_i\}_{i \in I}$ of $X_{\mathbb{C}}$ such that for each nonempty subset S of I , the set $V_S = \bigcap_{i \in S} V_i$ satisfies the following conditions.

- If nonempty, V_S is isomorphic to a simply connected domain in \mathbb{C} .
- The set $V_S \cap Z$ contains at most one element.

We then have

$$\chi_{\text{dR}}(X_{\mathbb{C}}^{\text{an}}, \mathcal{E}) = \sum_{S \subseteq I, S \neq \emptyset} (-1)^{\#S-1} \chi_{\text{dR}}(V_S, \mathcal{E}).$$

It then suffices to check that for each nonempty subset S of I ,

$$\chi_{\text{dR}}(V_S, \mathcal{E}) = \begin{cases} -\text{Irr}_z(\nabla) & (V_S \cap Z = \{z\}) \\ n & (V_S \cap Z = \emptyset, V_S \neq \emptyset) \\ 0 & (V_S = \emptyset). \end{cases} \quad (3)$$

In case $V_S = \emptyset$, there is nothing to check. In case $V_S \neq \emptyset$ but $V_S \cap Z = \emptyset$, this is immediate because the restriction of \mathcal{E} to V_S is trivial. In case $V_S \cap Z = \{z\}$, we may similarly replace V_S with a small open disc around z , and then invoke the Deligne–Malgrange interpretation of irregularity as the local index of meromorphic de Rham cohomology on a punctured disc [108, Théorème 3.3(d)]. \square

Theorem 7.7 (Global Index Formula). *We have*

$$\chi_{\text{dR}}(X^{\text{an}}, \mathcal{E}) = n\chi(U) - \int \Delta h(\mathcal{N}) = n(2 - 2g(X) - \text{length}(Z)) - \int \Delta h(\mathcal{N}). \quad (4)$$

Proof. This follows by comparing Lemma 7.4 with Lemma 7.6. \square

Remark 7.8. It may not be immediately obvious why Theorem 7.7 is of value, i.e., why it is useful to express the index of de Rham cohomology in terms of convergence polygons instead of irregularity. As observed by Baldassarri [11, (0.13)], there is a profound difference between the behavior of the index in the complex-analytic and nonarchimedean settings. In the complex case, for any open analytic subspace V of $X_{\mathbb{C}}^{\text{an}}$, we have

$$\chi_{\text{dR}}(V, \mathcal{E}) = n\chi(V \cap U^{\text{an}}) - \sum_{z \in V \cap Z} \text{Irr}_z(\nabla) \quad (5)$$

by the same argument as in the proof of (3). In particular, $\chi_{\text{dR}}(V, \mathcal{E}) \neq \chi_{\text{dR}}(V - Z, \mathcal{E})$; that is, the index of de Rham hypercohomology depends on whether we allow poles or essential singularities at the points of Z . By contrast, in the nonarchimedean case, these two indices coincide under a suitable technical hypothesis to ensure that they are both defined; see Example 7.9 for a simple example and Corollary 7.13 for the general case (and Example 7.19 and Remark 7.20 for a counterexample failing the technical hypothesis). This means that in the nonarchimedean case, the “source” of the index of de Rham cohomology is not irregularity, but rather the Laplacian of the convergence polygon (see Theorem 7.12).

Example 7.9. Consider the connection associated with Example 4.2 as per Remark 5.7; note that $\text{Irr}_0(\nabla) = 0$, $\text{Irr}_{\infty}(\nabla) = 1$, so $\chi_{\text{dR}}(X^{\text{an}}, \mathcal{E}) = -1$. For $0 < \alpha < \beta$, let V_{β}, W_{α} be the subspace $|z| < \beta, |z| > \alpha$ of X^{an} . By Mayer–Vietoris,

$$\chi_{\text{dR}}(V_{\beta}, \mathcal{E}) + \chi_{\text{dR}}(W_{\alpha}, \mathcal{E}) - \chi_{\text{dR}}(V_{\beta} \cap W_{\alpha}, \mathcal{E}) = \chi_{\text{dR}}(X^{\text{an}}, \mathcal{E}) = -1. \quad (6)$$

On the other hand, $\chi_{\mathrm{dR}}(V_\beta, \mathcal{E})$ (resp., $\chi_{\mathrm{dR}}(W_\alpha, \mathcal{E})$) equals the index of the operator $y \mapsto y' - y$ on Laurent series in z convergent for $|z| < \beta$ (resp., in z^{-1} convergent for $|z^{-1}| < \alpha^{-1}$). If $f = \sum_n f_n z^{-n}$, $g = \sum_n g_n z^{-n}$ are two such series, then the equation $g = f' - f$ is equivalent to

$$f_n = -g_n - (n-1)f_{n-1} \quad (n \in \mathbb{Z}). \quad (7)$$

Suppose first that $p^{-1/(p-1)} < \alpha < \beta$. Then given $g \in K((z^{-1}))$, we may solve uniquely for $f \in K((z^{-1}))$, and if g converges on $W_\alpha - \{\infty\}$, then so does f . We thus compute that

$$\chi_{\mathrm{dR}}(V_\beta, \mathcal{E}) = -1, \quad \chi_{\mathrm{dR}}(W_\alpha, \mathcal{E}) = 0, \quad \chi_{\mathrm{dR}}(V_\beta \cap W_\alpha, \mathcal{E}) = 0;$$

namely, the second equality is what we just computed, the third follows from Robba's index formula [146] (see also [97, Lemma 3.7.5]), and the first follows from the other two plus (6). In particular, the local index at ∞ equals 0, whereas in the complex-analytic setting it equals -1 by the Deligne–Malgrange formula (see the proof of Lemma 7.6).

Suppose next that $\alpha < \beta < p^{-1/(p-1)}$. Then by contrast, we have

$$\chi_{\mathrm{dR}}(V_\beta, \mathcal{E}) = 0, \quad \chi_{\mathrm{dR}}(W_\alpha, \mathcal{E}) = -1, \quad \chi_{\mathrm{dR}}(V_\beta \cap W_\alpha, \mathcal{E}) = 0;$$

namely, the first and third equalities follow from the triviality of ∇ on V_β , and the second follows from the other two plus (6).

With Example 7.9 in mind, we now describe a local refinement of Theorem 7.7, in which we dissect the combinatorial formula for the index into local contributions.

Definition 7.10. For $x \in U^{\mathrm{an}} \cap \Gamma_{X,Z}$, let $\mathrm{val}_\Gamma(x)$ be the valence of x as a vertex of $\Gamma_{X,Z}$, taking $\mathrm{val}_\Gamma(x) = 2$ when x lies on the interior of an edge. (We refer to *valence* instead of *degree* to avoid confusion with degrees of morphisms.) For $x \in U^{\mathrm{an}}$, define

$$\chi_x(\mathcal{E}) = \begin{cases} n(2 - 2g(C_x) - \mathrm{val}_\Gamma(x)) - (\Delta h(\mathcal{N}))_x & \text{if } x \in \Gamma_{X,Z} \\ -\Delta h(\mathcal{N})_x & \text{otherwise.} \end{cases} \quad (8)$$

Let $\chi(\mathcal{E})$ be the \mathbb{R} -valued measure whose value on a continuous function $f : U^{\mathrm{an}} \rightarrow \mathbb{R}$ is $\sum_{x \in U^{\mathrm{an}}} f(x) \chi_x(\mathcal{E})$.

Lemma 7.11. *We have*

$$\chi_{\mathrm{dR}}(X^{\mathrm{an}}, \mathcal{E}) = n\chi(U) - \int \Delta h(\mathcal{N}) = \int \chi(\mathcal{E}).$$

Proof. This follows from Theorem 7.7 plus the identity

$$\sum_{x \in U^{\text{an}} \cap \Gamma_{X,Z}} (2 - 2g(C_x) - \text{val}_\Gamma(x)) = 2 - 2g(X) - \text{length}(Z),$$

which amounts to the combinatorial formula for the genus of an analytic curve [9, Sect. 4.16]. \square

Theorem 7.12 (Local Index Formula). *Let V be an open subspace of X^{an} which is the retraction of an open subspace of $\Gamma_{X,Z}$. If $p > 0$, assume some additional technical hypotheses (see Remark 7.14). Then*

$$\chi_{\text{dR}}(V, \mathcal{E}) = \int_{V \cap U^{\text{an}}} \chi(\mathcal{E}).$$

Proof. See [135, Theorem 3.5.2]. \square

Corollary 7.13. *With hypotheses as in Theorem 7.12, $\chi_{\text{dR}}(V, \mathcal{E}) = \chi_{\text{dR}}(V \cap U^{\text{an}}, \mathcal{E})$; that is, the index of de Rham hypercohomology is the same whether we allow poles or essential singularities at Z .*

Remark 7.14. Let Γ be a strict skeleton for which the conclusion of Theorem 5.5 holds. For v a vertex of Γ , define the *star* of v , denoted \star_v , as the union of v and the interiors of the edges of Γ incident to v . Let $\pi_\Gamma : X \rightarrow \Gamma$ be the retraction onto Γ , through which \mathcal{N} factors. Under the hypotheses of Theorem 7.12, we have

$$\chi_{\text{dR}}(\pi_\Gamma^{-1}(\star_v), \mathcal{E}) = \chi_v(\mathcal{E}), \quad \chi_{\text{dR}}(\pi_\Gamma^{-1}(\star_v \cap \star_w), \mathcal{E}) = 0. \quad (9)$$

We can then recover Theorem 7.7 from (9) by using Mayer–Vietoris (and GAGA over K ; see Remark 7.8) to write

$$\chi_{\text{dR}}(X^{\text{an}}, \mathcal{E}) = \sum_v \chi_{\text{dR}}(\pi_\Gamma^{-1}(\star_v), \mathcal{E}) - \sum_{v \neq w} \chi_{\text{dR}}(\pi_\Gamma^{-1}(\star_v \cap \star_w), \mathcal{E});$$

one may similarly deduce Theorem 7.12 from (9).

Remark 7.15. For a given connection, one can extend the range of applicability of Theorem 7.12 by enlarging the set Z ; however, this depends on an understanding of how the two sides of the formula depend on Z . To this end, let us rewrite the equality as

$$\chi_{\text{dR}}(V, \mathcal{E}, Z) = \int_{V \cap U^{\text{an}}} \chi(\mathcal{E}, Z)$$

with Z included in the notation.

Put $Z' = Z \cup \{z'\}$ for some $z' \in U(K)$. Let $x' \in \Gamma_{X,Z}$ be the generic point of the open disc $U_{z'}$; by subdividing if necessary, we may view x' as a vertex of $\Gamma_{X,Z}$. Then $\Gamma_{X,Z'}$ is the union of $\Gamma_{X,Z}$ with a single edge joining x' to z' within $U_{z'}$.

Let V be an open subset of X^{an} containing z' which is the retraction of an open subspace of $\Gamma_{X,Z}$; it is then also a retraction of an open subspace of $\Gamma_{X,Z'}$. We then have

$$\chi_{\text{dR}}(V, \mathcal{E}, Z') - \chi_{\text{dR}}(V, \mathcal{E}, Z) = -n;$$

namely, by Mayer–Vietoris this reduces to the case where V is a small open disc around z' , in which case we may assume \mathcal{E} is trivial on V and make the computation directly.

By this computation, Theorem 7.12, and the fact that $\chi_{z'}(\mathcal{E}, Z) = 0$, we must also have

$$\int_{(V \cap U^{\text{an}}) \setminus \{z'\}} (\chi(\mathcal{E}, Z') - \chi(\mathcal{E}, Z)) = -n.$$

From (8), we see that $\chi(\mathcal{E}, Z') - \chi(\mathcal{E}, Z)$ is supported within $U_{z'} \cup \{x'\}$. However, while one can easily compute the convergence polygon associated with Z' from the one associated with Z (see Example 4.2 and Remark 8.13, for example), we do not know how to predict *a priori* where the support of $\chi(\mathcal{E}, Z') - \chi(\mathcal{E}, Z)$ will lie within $U_{z'} \cup \{x'\}$.

Remark 7.16. One of the main reasons we have restricted attention to meromorphic connections on proper curves is that in this setting, Theorem 5.5 ensures that $\chi_x(\mathcal{E}) = 0$ for all but finitely many $x \in U^{\text{an}}$. It is ultimately more natural to state Theorem 7.12 for connections on open analytic curves, as is done in [135, Theorem 3.8.10]; however, this requires some additional hypotheses to ensure that $\chi(\mathcal{E})$ is a finite measure.

Definition 7.17. Assume $p > 0$. A *p -adic Liouville number* is an element $x \in \mathbb{Z}_p - \mathbb{Z}$ such that

$$\liminf_{m \rightarrow \infty} \left\{ \frac{|y|}{m} : y \in \mathbb{Z}, y - x \in p^m \mathbb{Z}_p \right\} < +\infty.$$

As in the classical case, p -adic Liouville numbers are always transcendental [66, Proposition VI.1.1].

Remark 7.18. Assume $p > 0$. The technical hypotheses of Theorem 7.12 are needed to guarantee the existence of the indices appearing in (9). In case ∇ has a regular singularity at $z \in Z$ with all exponents in \mathbb{Z}_p , these hypotheses include the condition that no two exponents of ∇ at z differ by a p -adic Liouville number; see Example 7.19 for a demonstration of the necessity of such a condition. (Such hypotheses are not needed in Theorem 7.7 because there we only allow poles rather than essential singularities.)

Unfortunately, the full hypotheses are somewhat more complicated to state. They arise from the fact that with notation as in Remark 7.14, one can separate off a maximal component of \mathcal{E} on $\pi_{\Gamma}^{-1}(\star_v \cap \star_w)$ which satisfies the Robba condition, to which one may associate some p -adic numbers playing the role of exponents; the

hypothesis is that (for any particular v, w) no two of these numbers differ by p -adic Liouville numbers. The difficulty is that the definition of these p -adic exponents, due to Christol and Mebkhout (and later simplified by Dwork) is somewhat indirect; they occur as “resonant frequencies” for a certain action by the group of p -power roots of unity, which are hard to control except in some isolated cases where they are forced to be rational numbers (e.g., Picard–Fuchs equations, a/k -Gauss–Manin connections, or connections arising from F -isocrystals in the theory of crystalline cohomology). See [92, Chap. 13] for more discussion.

Example 7.19. Assume $p > 0$ and take $X = \mathbb{P}_K^1, Z = \{0, \infty\}$. Take \mathcal{E} to be free of rank 1 with the action of ∇ given by

$$\nabla(f) = \lambda f \frac{dz}{z} + df$$

for some $\lambda \in K$. For α, β with $0 < \alpha < \beta$, let V be the open annulus $\alpha < |z| < \beta$. The 1-forms on V are series $\sum_{n=-\infty}^{\infty} c_n z^n \frac{dz}{z}$ such that for each $\rho \in (\alpha, \beta)$, $|c_n| \rho^n \rightarrow 0$ as $n \rightarrow \pm\infty$.

If $\lambda = 0$, then $|c_n| \rho^n \rightarrow 0$ for all $\rho \in (\alpha, \beta)$ if and only if $|c_n/n| \rho^n \rightarrow 0$ for all $\rho \in (\alpha, \beta)$, so every 1-form on V with $c_0 = 0$ is in the image of ∇ . Note that multiplying the generator of \mathcal{E} by t has the effect of replacing λ by $\lambda + 1$; it follows that if $\lambda \in \mathbb{Z}$, then the kernel and cokernel of ∇ on V are both 1-dimensional, so $\chi_{\text{dR}}(V, \mathcal{E}) = 0$.

If $\lambda \in K - \mathbb{Z}$, then a 1-form is in the image of ∇ on V if and only if for each $\rho \in (\alpha, \beta)$, $|c_n/(n-\lambda)| \rho^n \rightarrow 0$ as $n \rightarrow \pm\infty$. This holds if λ is not a p -adic Liouville number (see [66, Sect. VI.1] or [92, Proposition 13.1.4]); otherwise, one shows that ∇ has infinite-dimensional cokernel on V , so $\chi_{\text{dR}}(V, \mathcal{E})$ is undefined.

Remark 7.20. Example 7.19 provides an example showing that the equality $\chi_{\text{dR}}(V, \mathcal{E}) = \chi_{\text{dR}}(V \cap U^{\text{an}}, \mathcal{E})$ of Corollary 7.13 cannot hold without conditions on p -adic Liouville numbers. In this example, for all λ , $\chi_{\text{dR}}(X^{\text{an}}, \mathcal{E}) = 0$ by Lemma 7.6; but, when λ is a p -adic Liouville number, $\chi_{\text{dR}}(U^{\text{an}}, \mathcal{E})$ is undefined.

Remark 7.21. The net result of Remark 7.18 is that in general, one can only view $\chi_x(\mathcal{E})$ as a *virtual local index* of \mathcal{E} at x , not a true local index. Nonetheless, this interpretation can be used to predict combinatorial properties of the convergence polygon which often continue to hold even without restrictions on p -adic exponents. For example, Theorem 6.6 corresponds to the fact that if V is an open disc in U^{an} , then the dimension of the cokernel of ∇ on V is nonnegative; this argument appears in the proof of Theorem 6.6 (see Theorem A.6).

Remark 7.22. In light of Remark 7.21, one might hope to establish some inequalities on $\chi_x(\mathcal{E})$. One might first hope to refine Theorem 7.23 by analogy with (5), by proving that the measure $\Delta h(\mathcal{N})$ is nonnegative; however, this fails already in simple examples such as Example 7.25.

On the other hand, since our running hypothesis is that $\chi(U) \leq 0$, Theorem 7.7 and Lemma 7.11 imply that

$$\sum_{x \in U^{\text{an}}} \chi_x(\mathcal{E}) = \chi_{\text{dR}}(X^{\text{an}}, \mathcal{E}) \leq n\chi(U) \leq 0.$$

One might thus hope to refine Theorem 7.23 by proving that $\chi_x(\mathcal{E}) \leq 0$ for all $x \in U^{\text{an}}$. Unfortunately, this is not known (and may not even be safe to conjecture) in full generality, but see Theorem 7.23 for some important special cases.

Theorem 7.23. *Choose $x \in U^{\text{an}}$.*

- (a) *Let $R(x)$ be the infimum of the radii of open discs in U_x containing x . If $\mathcal{N}(x)$ has no slopes equal to $-\log R(x)$, then $\chi_x(\mathcal{E}) = 0$.*
- (b) *If $x \in \Gamma_{X,Z}$, then $\chi_x(\mathcal{E}) \leq 0$.*
- (c) *If $p = 0$, then $\chi_x(\mathcal{E}) \leq 0$.*

Proof. For (a), see [133, Proposition 6.2.11] or [16] (or reduce to the case where $\mathcal{N}(x)$ has all slopes greater than $-\log R(x)$ and then apply [97, Theorem 5.3.6]). A similar argument implies (b) because in this case, $R(x) = 1$ and the zero slopes are forced to make a nonpositive contribution to the index. For (c), see Theorem A.9. \square

Remark 7.24. The proof of Theorem 7.23 can also be used to quantify the extent to which \mathcal{N} fails to factor through the retract onto $\Gamma_{X,Z}$, and hence to help identify a suitable skeleton Γ for which the conclusion of Theorem 5.5 holds. To be precise, for $x \in \Gamma_{X,Z}$, if the restriction of \mathcal{N} to $\Gamma_{X,Z}$ is harmonic at x , then \mathcal{N} is constant on the fiber at x of the retraction of X^{an} onto $\Gamma_{X,Z}$. See also [133, Sect. 6.3].

Example 7.25. Let h be a nonzero rational function on X , and take Z to be the pole locus of h . Let \mathcal{E} be the free bundle on a single generator \mathbf{v} equipped with the connection

$$\nabla(f\mathbf{v}) = \mathbf{v} \otimes (df + f dh).$$

For each $z \in Z$, $\text{Irr}_z(\nabla)$ equals the multiplicity of z as a pole of h . By Lemma 7.11,

$$\sum_{x \in U^{\text{an}}} \chi_x(\mathcal{E}) = \chi(U) - m$$

where m is the number of poles of h counted with multiplicity.

Suppose now that there exists a point $x \in U^{\text{an}}$ with $g(C_x) > 0$. By multiplying h by a suitably large element of K , we can ensure that $\mathcal{N}(x)$ has positive slope. In this case, by Theorem 7.23 we must have

$$\chi_x(\mathcal{E}) = 0 > 2 - 2g(C_x) - \text{val}_\Gamma(x).$$

In particular, while $\int \Delta h(\mathcal{N}) = \sum_{z \in Z} \text{Irr}_z(\nabla) \geq 0$, the measure $\Delta h(\mathcal{N})$ is not necessarily nonnegative.

8 Ramification of Finite Morphisms

Hypothesis 8.1. Throughout Sect. 8, let $f : Y \rightarrow X$ be a finite flat morphism of degree d of smooth, proper, geometrically connected curves over K (a non-archimedean field of characteristic 0), and suppose that the restriction of f to U (the complement of a finite set Z of closed points of X) is étale and Galois with Galois group G .

Let $\pi : G \rightarrow \mathrm{GL}(V)$ be a faithful representation of G on an n -dimensional vector space V over K . Note that since G is finite and π is faithful (and K is of characteristic 0), the Tannakian category generated by V contains all finite-dimensional K -linear representations of G , including the regular representation.

Definition 8.2. Equip $\mathcal{E}_V = \mathcal{O}_Y \otimes_K V^\vee$ with the diagonal action of G induced by the Galois action on \mathcal{O}_Y and the action via π^\vee on V^\vee . By faithfully flat descent, G -invariant sections of \mathcal{E}_V can be identified with sections of a vector bundle \mathcal{E} on U ; moreover, the trivial connection on \mathcal{E}_V induces a connection ∇ on \mathcal{E} . We are thus in the situation of Hypothesis 5.1.

Remark 8.3. One way to arrive at Hypothesis 8.1 is to start with a non-Galois cover $g : W \rightarrow X$, let f be the Galois closure, and let π be the representation of G induced by the trivial representation of the subgroup of G fixing W . In this case, \mathcal{E} is just the pushforward of the trivial connection on $g^{-1}(U)$.

Remark 8.4. Following up on the previous remark, note that (\mathcal{E}, ∇) is always a subobject of the pushforward of the trivial connection on $f^{-1}(U)$. Since the latter is the Gauss–Manin connection associated with a smooth proper morphism (of relative dimension 0), it is everywhere regular by virtue of the Griffiths monodromy theorem [82, Theorem 14.1]. This implies in turn that for $z \in Z$, our original connection satisfies $\mathrm{Irr}_z(\nabla) = 0$ and the exponents of ∇ at z are rational.

To say more, let N be the order of the inertia subgroup of G corresponding to some element of $f^{-1}(\{z\})$; we then claim that the exponents of ∇ at z belong to $\frac{1}{N}\mathbb{Z}$ and generate this abelian group. To check this, we may assume that K is algebraically closed. By the Cohen structure theorem, the field $\mathrm{Frac}(\widehat{\mathcal{O}}_{X,z})$ may be identified with $K((t))$. The vector space $\mathcal{E}_z = \mathcal{E} \otimes_{\mathcal{O}_U} \mathrm{Frac}(\widehat{\mathcal{O}}_{X,z})$ over this field splits compatibly with ∇ as a direct sum

$$\bigoplus_{y \in f^{-1}(\{z\})} \mathrm{Frac}(\widehat{\mathcal{O}}_{Y,y}),$$

each summand of which is isomorphic to $K((t^{1/N}))$. If we split this summand further as $\bigoplus_{i=0}^{N-1} t^{i/N} K((t))$, then the i -th summand contributes an exponent congruent to i/N modulo \mathbb{Z} . In particular, when π is the regular representation, the exponents of ∇ at z belong to $\frac{1}{N}\mathbb{Z}$ and fill out the quotient $\frac{1}{N}\mathbb{Z}/\mathbb{Z}$. For a general π , on one hand, π occurs as a summand of the regular representation, so the exponents of ∇ at z again belong to $\frac{1}{N}\mathbb{Z}$. On the other hand, the regular representation occurs as a

summand of some tensor product of copies of V and its dual (see Hypothesis 8.1), so the group generated by the exponents must also fill out the quotient $\frac{1}{N}\mathbb{Z}/\mathbb{Z}$ (although the exponents themselves need not).

Definition 8.5. For L a nonarchimedean field over K , let \mathbb{C}_L denote a completed algebraic closure of L , and let B_L be the subset of X_L for which $x \in B_L$ if and only if the preimage of x in $Y_{\mathbb{C}_L}^{\text{an}}$ does not consist of d distinct points. (For a demonstration of this definition, see Example 8.12.)

We have the following remark about B_L echoing an earlier observation about $\Gamma_{X,Z}$ (see Definition 2.5).

Remark 8.6. Let L' be a nonarchimedean field containing L . Then $B_{L'}$ is contained in the inverse image of B_L under the restriction map $X_{L'}^{\text{an}} \rightarrow X_L^{\text{an}}$, but this containment is typically strict if L' is not the completion of an algebraic extension of L . This is because $B_{L'}$ only contains points of types 2 or 3 (except for the points of Z), whereas the inverse image of B_L typically also contains some points of type 1.

Example 8.7. Take $K = \mathbb{Q}_p(\zeta_p)$, $X = \mathbb{P}_K^1$, $Z = \{0, \infty\}$, $Y = \mathbb{P}_K^1$, let $f : Y \rightarrow X$ be the map $z \mapsto z^p$, identify G with $\mathbb{Z}/p\mathbb{Z}$ so that $1 \in \mathbb{Z}/p\mathbb{Z}$ corresponds to the map $z \mapsto \zeta z$ on $f^{-1}(U)$, and let $\pi : G \rightarrow \text{GL}_1(K)$ be the character taking 1 to ζ_p^{-1} . Then \mathcal{E} is free on a single generator \mathbf{v} satisfying $\nabla(\mathbf{v}) = -\frac{1}{p}z^{-1}\mathbf{v} \otimes dz$. This is also the connection obtained from Example 4.3 by applying Remark 5.7.

Put $\omega = p^{-1/(p-1)} = |\zeta_p - 1|$. For each $x \in U$, we may define the *normalized diameter* of x as an element of $[0, 1]$ defined as follows: choose an extension L of K such that x lifts to some $\tilde{x} \in U(L)$, then take the infimum of all $\rho \in (0, 1]$ such that $U_{\tilde{x}, \rho}$ meets the inverse image of x in U_L^{an} (or 1 if no such ρ exists). With this definition, for any L , the set B_L consists of Z plus all points with normalized diameter in $[\omega^p, 1]$ (see, for example, [92, Lemma 10.2.2]).

Lemma 8.8. *Suppose that L is algebraically closed. Let V be an open subset of X_L^{an} which is disjoint from B_L . Then the map $f^{-1}(V) \rightarrow V$ of topological spaces is a covering space map.*

Proof. For each $x \in V$, the local ring of X_L^{an} at x is Henselian [17, Theorem 2.1.5]; hence, for each $y \in f^{-1}(\{x\})$, we can find a neighborhood W_y of y in $f^{-1}(V)$ which is finite étale over its image in X_L^{an} . Since Y_L^{an} is Hausdorff, we can shrink the W_y so that they are pairwise disjoint and all have the same image V' in X_L^{an} . The maps $W_y \rightarrow V'$ are then finite étale of some degrees adding up to $\deg(f) = d$. Since $\#f^{-1}(\{x\}) = d$ by hypothesis, these degrees must all be equal to 1; hence, each map $W_y \rightarrow V'$ is actually an isomorphism of analytic spaces, and in particular a homeomorphism of topological spaces. \square

Theorem 8.9. *For L a nonarchimedean field over K and $x \in U(L)$, $e^{-s_1(\mathcal{N}(x))}$ equals the supremum of $\rho \in (0, 1]$ such that $U_{x, \rho} \cap B_L = \emptyset$.*

Proof. We may assume without loss of generality that L is itself algebraically closed. If $U_{x, \rho} \cap B_L = \emptyset$, then by Lemma 8.8, the map $f^{-1}(U_{x, \rho}) \rightarrow U_{x, \rho}$ is

a covering space map of topological spaces. Since $U_{x,\rho}$ is contractible and hence simply connected, $f^{-1}(U_{x,\rho})$ splits topologically as a disjoint union of copies of $U_{x,\rho}$. From this, it follows easily that the restriction of \mathcal{E} to $U_{x,\rho}$ is trivial.

Conversely, suppose that the restriction of \mathcal{E} to $U_{x,\rho}$ is trivial. Then the same holds when V is replaced by any representation in the Tannakian category generated by V ; since G is finite and π is faithful, this includes the regular representation. That is, the pushforward of the trivial connection on $f^{-1}(U)$ has trivial restriction to $U_{x,\rho}$.

In addition to the connection, the sections of this restriction inherit from \mathcal{O}_Y a multiplication map, and the product of two horizontal sections is again horizontal. Consequently, the horizontal sections form a finite reduced L -algebra of rank d ; since L is algebraically closed, this algebra must split as a direct sum of d copies of L . This splitting corresponds to a splitting of $f^{-1}(U_{x,\rho})$ into d disjoint sets, proving that $U_{x,\rho} \cap B_L = \emptyset$. □

Remark 8.10. One may view Theorem 8.9 as saying that under a suitable normalization, $e^{-s_1(\mathcal{N}(x))}$ measures the distance from x to B_L . This suggests the interpretation of B_L as an “extended ramification locus” of the map f ; for maps from \mathbb{P}_K^1 to itself, this interpretation has been adopted in the context of nonarchimedean dynamics (e.g., see [67, 68]). However, this picture is complicated by Remark 8.6; roughly speaking, even for $x \in B_L$, the “distance from x to itself” must be interpreted as a nonzero quantity. In any case, Theorem 8.9 suggests the possibility of relating the full convergence polygon to more subtle measures of ramification, such as those considered recently by Temkin [47, 157]. One older result in this direction is the theorem of Matsuda and Tsuzuki; see Theorem 8.11 below.

Theorem 8.11. *Assume that K has perfect residue field of characteristic $p > 0$. Suppose that $x \in \Gamma_{X,Z}$ is of type 2, choose $y \in f^{-1}(\{x\})$, and suppose that $\mathcal{H}(y)$ is unramified over $\mathcal{H}(x)$. Let κ_x, κ_y be the residue fields of the nonarchimedean fields $\mathcal{H}(x), \mathcal{H}(y)$. Let $\bar{\pi}$ be the restriction of π along the identification of $H := \text{Gal}(\kappa_y/\kappa_x) \cong \text{Gal}(\mathcal{H}(y)/\mathcal{H}(x))$ with the stabilizer of y in G .*

- (a) *The polygon $\mathcal{N}(x)$ is zero.*
- (b) *Let \bar{t} be a branch of X at x , and let v be the point on C_x corresponding to \bar{t} (see Definition 6.1). Then $\partial_{\bar{t}}(\mathcal{N})$ computes the wild Hasse–Arf polygon of $\bar{\pi}$ at v , i.e., the Newton polygon in which the slope $s \geq 0$ occurs with multiplicity $\dim_K(V^{H^s+}/V^{H^s})$. In particular, $\partial_{\bar{t}}(h(\mathcal{N}))$ computes the Swan conductor of $\bar{\pi}$ at v .*

Proof. See any of [87, 114, 163]. □

Example 8.12. Set notation as in Example 8.7, except take $Z = \{0, 1, \infty\}$. Let $x \in \Gamma_{X,Z}$ be the generic point of the disc $|z - 1| \leq \omega^p$, which we can also write as $|u| \leq 1$ for $u = (z - 1)/(\zeta_p - 1)^p$. Let y be the unique preimage of x in Y . Let t be the coordinate $(z - 1)/(\zeta_p - 1)$ on Y ; then, $\kappa_x = \mathbb{F}_p(\bar{u})$ while $\kappa_y = \mathbb{F}_p(\bar{t})$ where $\bar{t}^p - \bar{t} = \bar{u}$. For \bar{t} the branch of x towards x_1 , we have $\partial_{\bar{t}}(h(\mathcal{N})) = 1$; as predicted by Theorem 8.11, this equals the Swan conductor of the residual extension at ∞ .

Remark 8.13. In Example 8.12, it is necessary to include 1 in Z to force x into $\Gamma_{X,Z}$, so that Theorem 8.11 applies; otherwise, we would have $\partial_{\bar{t}}(h(\mathcal{N})) = 0$ as in Example 4.3. This provides an explicit example of the effect of enlarging Z on the behavior of \mathcal{N} , as described in Remark 7.15.

Remark 8.14. To generalize Theorem 8.11 to cases where $\mathcal{H}(y)$ is ramified over $\mathcal{H}(x)$, it may be most convenient to use Huber’s ramification theory for adic curves [79], possibly as refined in [47, 157]. In a similar vein, the global index formula (Theorem 7.7) is essentially the Riemann–Hurwitz formula for the map f , in which case it should be possible to match up the local contributions appearing in Theorem 7.7 with ramification-theoretic local contributions.

Remark 8.15. Suppose that $p > 0$, $X = \mathbb{P}_K^1$, and f extends to a finite flat morphism of smooth curves over \mathfrak{o}_K with target $\mathbb{P}_{\mathfrak{o}_K}^1$. By [51, Theorem 3.1], \mathcal{E} admits a unit-root Frobenius structure in a neighborhood of x_1 , from which it follows that \mathcal{E} satisfies the Robba condition at x_1 .

Let k be the residue field of K . If in addition the points of Z have distinct projections to \mathbb{P}_k^1 , then by Remark 5.11, \mathcal{E} satisfies the Robba condition everywhere. By Remark 6.4, this implies that \mathcal{E} has exponents in \mathbb{Z}_p ; by Remark 8.4, this implies that f is tamely ramified (i.e., the inertia group of each point of Y has order coprime to p). By contrast, if the points of Z do not have distinct projections to \mathbb{P}_k^1 , then f need not be tamely ramified; see Sect. 11.

Remark 8.16. For a connection derived from a finite morphism, in case $p > 0$ the technical conditions of Remark 7.18 are always satisfied. An important consequence of this statement for our present work is that the conclusion of Theorem 7.23(c) can be established for such a connection even when $p > 0$; see Remark A.11.

9 Artin–Hasse Exponentials and Witt Vectors

We will conclude by specializing the previous discussion to cyclic covers of discs in connection with the Oort local lifting problem. In preparation for this, we need to recall some standard constructions of p -adic analysis.

Hypothesis 9.1. Throughout Sect. 9, fix a prime p and a positive integer n . In some algebraic closure of \mathbb{Q}_p , fix a sequence of primitive p^n -th roots of unity ζ_{p^n} such that $\zeta_{p^n}^p = \zeta_{p^{n-1}}$.

Definition 9.2. The *Artin–Hasse exponential series* at p is the formal power series

$$E_p(t) := \exp \left(\sum_{i=0}^{\infty} \frac{t^{p^i}}{p^i} \right).$$

Lemma 9.3. We have $E_p(t) \in \mathbb{Z}_{(p)}[[t]]$. In particular, $E_p(t)$ converges for $|t| < 1$.

Proof. See, for instance, [92, Proposition 9.9.2]. □

Lemma 9.4. *Let $\mathbb{Z}_{(p)}\langle t \rangle$ be the subring of $\mathbb{Z}_{(p)}[[t]]$ consisting of series $\sum_{i=0}^{\infty} c_i t^i$ for which the c_i converge p -adically to 0 as $i \rightarrow \infty$. Then if we define the power series*

$$f(z, t) := \frac{E_p(zt)E_p(t^p)}{E_p(t)E_p(zt^p)} \in \mathbb{Z}_{(p)}[[t, 1-z]]$$

using Lemma 9.3, we have

$$f(z, t) \in z + (t-1)\mathbb{Z}_{(p)}\langle t \rangle[[1-z]].$$

Proof. Write $f(z, t) = \exp g(z, t)$ with

$$\begin{aligned} g(z, t) &= \sum_{i=0}^{\infty} \frac{(z^{p^i} - 1)(t^{p^i} - t^{p^{i+1}})}{p^i} \\ &= \sum_{i=0}^{\infty} \sum_{j=1}^{\infty} (-1)^j \binom{p^i}{j} (1-z)^j \frac{t^{p^i} - t^{p^{i+1}}}{p^i} \\ &= \sum_{j=1}^{\infty} \frac{(-1)^j}{j} (1-z)^j \sum_{i=0}^{\infty} \binom{p^i - 1}{j-1} (t^{p^i} - t^{p^{i+1}}) \\ &= \sum_{j=1}^{\infty} \frac{(-1)^j}{j} (1-z)^j \left(\binom{0}{j-1} t + \sum_{i=1}^{\infty} \left(\binom{p^i - 1}{j-1} - \binom{p^{i-1} - 1}{j-1} \right) t^{p^i} \right). \end{aligned}$$

For each fixed j , $\binom{p^i - 1}{j-1}$ converges p -adically to $\binom{-1}{j-1} = (-1)^{j-1}$ as $i \rightarrow \infty$. It follows that $g(z, t) \in (\mathbb{Z}_{(p)}\langle t \rangle \otimes_{\mathbb{Z}} \mathbb{Q})[[1-z]]$ and

$$g(z, t) \equiv - \sum_{j=1}^{\infty} \frac{(1-z)^j}{j} \pmod{(t-1)(\mathbb{Z}_{(p)}\langle t \rangle \otimes_{\mathbb{Z}} \mathbb{Q})[[1-z]]}.$$

This then implies that $f(z, t) \in (\mathbb{Z}_{(p)}\langle t \rangle \otimes_{\mathbb{Z}} \mathbb{Q})[[1-z]]$ and

$$f(z, t) \equiv z \pmod{(1-t)(\mathbb{Z}_{(p)}\langle t \rangle \otimes_{\mathbb{Z}} \mathbb{Q})[[1-z]]}.$$

Since $f(z, t)$ also belongs to $\mathbb{Z}_{(p)}[[t, 1-z]]$ by Lemma 9.3, we may deduce the claimed inclusion. \square

Definition 9.5. Let W_n denote the p -typical Witt vector functor. Given a ring R , the set $W_n(R)$ consists of n -tuples $\underline{a} = (a_0, \dots, a_{n-1})$, and the arithmetic operations on $W_n(R)$ are characterized by functoriality in R and the property that the ghost map $w : W_n(R) \rightarrow R^n$ given by

$$(a_0, \dots, a_{n-1}) \mapsto (w_0, \dots, w_{n-1}), \quad w_i = \sum_{j=0}^i p^j a_j^{p^{i-j}} \tag{10}$$

is a ring homomorphism for the product ring structure on R^n . For any ideal I of R , let $W_n(I)$ denote the subset of $W_n(R)$ consisting of n -tuples with components in I ; since $W_n(I) = \ker(W_n(R) \rightarrow W_n(R/I))$, it is an ideal of $W_n(R)$.

Various standard properties of Witt vectors may be derived by using functoriality to reduce to polynomial identities over \mathbb{Z} , then checking these over \mathbb{Q} using the fact that the ghost map is a bijection if $p^{-1} \in R$. Here are two key examples.

- (i) Define the *Teichmüller map* sending $r \in R$ to $[r] := (r, 0, \dots, 0) \in W_n(R)$. Then this map is multiplicative: for all $r, s \in R$, $[rs] = [r][s]$.
- (ii) Define the *Verschiebung map* sending $\underline{a} \in W_n(R)$ to $V(\underline{a}) := (0, a_0, \dots, a_{n-2}) \in W_n(R)$. Then this map is additive: for all $\underline{a}, \underline{b} \in W_n(R)$, $V_n(\underline{a} + \underline{b}) = V_n(\underline{a}) + V_n(\underline{b})$.

Definition 9.6. In case R is an \mathbb{F}_p -algebra, the Frobenius endomorphism $\varphi : R \rightarrow R$ extends by functoriality to $W_n(R)$ and satisfies

$$p\underline{a} = (V \circ \varphi)(\underline{a}) = (\varphi \circ V)(\underline{a}) \quad (\underline{a} \in W_n(R)); \quad (11)$$

see, for instance, [80, Sect. 0.1]. It follows that for general R , if we define the map σ sending $\underline{a} \in W_n(R)$ to $\sigma(\underline{a}) = (a_0^p, \dots, a_{n-1}^p) \in W_n(R)$, then

$$p\underline{a} - (V \circ \sigma)(\underline{a}) \in W_n(pR) \quad (\underline{a} \in W_n(R)). \quad (12)$$

Beware that σ is in general not a ring homomorphism.

Definition 9.7. By Lemma 9.3, we have

$$E_{n,p}(t) := \frac{E_p(\zeta_{p^n} t)}{E_p(t)} = \exp \left(\sum_{i=0}^{n-1} (\zeta_{p^{n-i}} - 1) \frac{t^{p^i}}{p^i} \right) \in \mathbb{Z}_{(p)}[\zeta_{p^n}][[t]].$$

We may also define

$$E_{n,p}(\underline{a}) := \prod_{i=0}^{n-1} E_{n-i,p}(a_i) \in \mathbb{Z}_{(p)}[\zeta_{p^n}][[a_0, \dots, a_{n-1}]],$$

which we may also write as

$$E_{n,p}(\underline{a}) = \exp \left(\sum_{i=0}^{n-1} (\zeta_{p^{n-i}} - 1) \frac{w_i}{p^i} \right) \quad (13)$$

for w_i as in (10). Consequently, in $\mathbb{Z}_{(p)}[\zeta_{p^n}][[a_0, \dots, a_{n-1}, b_0, \dots, b_{n-1}]]$,

$$E_{n,p}(\underline{a})E_{n,p}(\underline{b}) = E_{n,p}(\underline{a} + \underline{b}). \quad (14)$$

Definition 9.8. By Lemma 9.3, we may define the formal power series

$$F_{n,p}(t) := \frac{E_{n,p}(t)}{E_{n,p}(t^p)} = \exp \left(\sum_{i=0}^{n-1} \frac{(\zeta_{p^{n-i}} - 1)(t^{p^i} - t^{p^{i+1}})}{p^i} \right) \in \mathbb{Z}_{(p)}[\zeta_{p^n}][[t]]$$

and

$$F_{n,p}(\underline{a}) := \frac{E_{n,p}(\underline{a})}{E_{n,p}(\sigma(\underline{a}))} = \prod_{i=0}^{n-1} F_{n-i,p}(a_i) \in \mathbb{Z}_{(p)}[\zeta_{p^n}][[a_0, \dots, a_{n-1}]].$$

By (13), in $\mathbb{Z}_{(p)}[\zeta_{p^n}][[a_0, \dots, a_{n-1}, b_0, \dots, b_{n-1}]]$ we have

$$F_{n,p}(\underline{a})F_{n,p}(\underline{b}) = F_{n,p}(\underline{a} + \underline{b})E_{n,p}(\sigma(\underline{a} + \underline{b}) - \sigma(\underline{a}) - \sigma(\underline{b})). \quad (15)$$

We will see shortly (Lemma 9.10) that $F_{n,p}(t)$ has radius of convergence greater than 1, which implies an analogous assertion for $F_{n,p}(\underline{a})$. For $p > 2$, this is shown in [113, Proposition 1.10] using a detailed computational argument; our argument follows the more conceptual approach given in [139, Theorem 2.5].

Definition 9.9. By Lemma 9.3, we may define the formal power series

$$G_{n,p}(\underline{a}) = \frac{E_{n,p}(p\underline{a})}{E_{n-1,p}(\underline{a})} \in \mathbb{Z}_{(p)}[\zeta_{p^n}][[a_0, \dots, a_{n-1}]].$$

By (14),

$$G_{n,p}(\underline{a})G_{n,p}(\underline{b}) = G_{n,p}(\underline{a} + \underline{b}). \quad (16)$$

We also have

$$G_{n,p}(\underline{a}) = E_{n,p}(p\underline{a} - (0, a_0^p, \dots, a_{n-2}^p))F_{n-1,p}(\underline{a})^{-1}. \quad (17)$$

Lemma 9.10. (a) *We have*

$$G_{n,p}(\underline{a}) \in \mathbb{Z}_{(p)}[\zeta_{p^n}][[(\zeta_{p^n} - 1)a_0, \dots, (\zeta_p - 1)a_{n-1}]].$$

(b) *The power series $F_{n,p}(\underline{a})$ converges on a polydisc with radius of convergence strictly greater than 1.*

Proof. Write

$$\begin{aligned} G_{n,p}(\underline{a}) &= \prod_{j=0}^{n-1} \exp \left(\sum_{i=0}^{n-j-1} (p(\zeta_{p^{n-j-i}} - 1) - (\zeta_{p^{n-j-1-i}} - 1)) \frac{a_j^{p^i}}{p^i} \right) \\ &= \prod_{j=0}^{n-1} \exp \left(\sum_{i=0}^{n-j-1} (\zeta_{p^{n-j-i}} - 1)(p - 1 - \zeta_{p^{n-j-i}} - \dots - \zeta_{p^{n-j-i}}^{p-1}) \frac{a_j^{p^i}}{p^i} \right) \\ &= \prod_{j=0}^{n-1} E_{n-j,p} \left((p - 1 - [\zeta_{p^{n-j}}] - \dots - [\zeta_{p^{n-j}}^{p-1}])[a_j] \right). \end{aligned}$$

Note that $p - 1 - [\zeta_{p^{n-j}}] - \dots - [\zeta_{p^{n-j}}^{p-1}]$ maps to zero in $W_{n-j}(\mathbb{F}_p)$ and so belongs to $W_{n-j}((\zeta_{p^{n-j}} - 1)\mathbb{Z}_{(p)}[\zeta_{p^{n-j}}])$. This proves (a).

To prove (b), apply (17) to write

$$F_{n,p}(\underline{a}) = G_{n+1,p}(\underline{a})^{-1} E_{n+1,p}(p\underline{a} - (V \circ \sigma)(\underline{a})).$$

Since $G_{n+1,p}(0) = 1$, (a) implies that $G_{n+1,p}(\underline{a})^{-1}$ converges on a polydisc with radius of convergence strictly greater than 1. Since $E_{n+1,p}(\underline{a}) \in \mathbb{Z}_{(p)}[\zeta_{p^{n+1}}][[a_0, \dots, a_{n-1}]]$ and (12) implies that $p\underline{a} - (V \circ \sigma)(\underline{a}) \in W_n(pR)$, $E_{n+1,p}(p\underline{a} - (V \circ \sigma)(\underline{a}))$ also converges on a polydisc with radius of convergence strictly greater than 1. This proves (b). \square

Lemma 9.11. *For $m \in \mathbb{Z}$, we have*

$$F_{n,p}(\underline{a} + \underline{b} + m) F_{n,p}(\underline{a})^{-1} E_{n,p}(\sigma(\underline{a} + \underline{b} + m) - \sigma(\underline{a}) - \sigma(\underline{b}) - m) = \zeta_{p^n}^m \quad (18)$$

as an equality of elements of $\mathbb{Z}_p[\zeta_{p^n}][[a_0, \dots, a_{n-1}, b_0, \dots, b_{n-1}]]$.

Note that the presence of m prevents us from heedlessly applying (14), because, for instance, $E_{n,p}(m)$ does not make sense. (We like to think of this as an example of *conditional convergence* in a nonarchimedean setting.) Similarly, we must work over \mathbb{Z}_p rather than $\mathbb{Z}_{(p)}$.

Proof. Convergence of $F_{n,p}(\underline{a} + \underline{b} + m)$ is guaranteed by Lemma 9.10. Convergence of $E_{n,p}(\sigma(\underline{a} + \underline{b} + m) - \sigma(\underline{a}) - \sigma(\underline{b}) - m)$ is guaranteed by Lemma 9.3 and the fact that

$$\sigma(\underline{a} + \underline{b} + m) - \sigma(\underline{a}) - \sigma(\underline{b}) - m \in W_n((\zeta_{p^n} - 1)\mathbb{Z}_p[\zeta_{p^n}][[a_0, \dots, a_{n-1}, b_0, \dots, b_{n-1}]]).$$

Thus the left side of (18) is well-defined. Using (14), we see that this quantity is constant as a power series in $a_0, \dots, a_{n-1}, b_0, \dots, b_{n-1}$; it thus remains to prove that

$$F_{n,p}(m) E_{n,p}(\sigma(m) - m) = \zeta_{p^n}^m. \quad (19)$$

Using (14) and (15), we see that for $m, m' \in \mathbb{Z}$,

$$\begin{aligned} & F_{n,p}(m) E_{n,p}(\sigma(m) - m) F_{n,p}(m') E_{n,p}(\sigma(m') - m') \\ &= F_{n,p}(m + m') E_{n,p}(\sigma(m + m') - \sigma(m) - \sigma(m')) \\ & \quad E_{n,p}(\sigma(m) - m) E_{n,p}(\sigma(m') - m') \\ &= F_{n,p}(m + m') E_{n,p}(\sigma(m + m') - m - m'), \end{aligned}$$

so both sides of (19) are multiplicative in m . It thus suffices to check (19) for $m = 1 = (1, 0, \dots)$, in which case $\sigma(m) = m$ and so the desired equality becomes $F_{n,p}(1) = \zeta_{p^n}$. This follows from Lemma 9.4 by evaluating at $z = \zeta_{p^n}$. \square

10 Kummer–Artin–Schreier–Witt Theory

In further preparation for discussion of the Oort local lifting problem, we describe a form of *Kummer–Artin–Schreier–Witt theory* for cyclic Galois extensions of a power series field.

Hypothesis 10.1. Throughout Sect. 10, fix a positive integer n , assume that K is discretely valued, the residue field k of K is algebraically closed of characteristic $p > 0$, and K contains a primitive p^n -th root of unity ζ_{p^n} . Put $F = k(\overline{\zeta})$.

Definition 10.2. For $\rho \in (0, 1)$, let $A(\rho, 1)$ be the annulus $\rho < |z| < 1$ in \mathbb{P}_K^1 ; the analytic functions on $A(\rho, 1)$ can be viewed as certain Laurent series in z . The union of the rings $\mathcal{O}(A(\rho, 1))$ over all $\rho \in (0, 1)$ is called the *Robba ring* over K and will be denoted \mathcal{R} . (This ring can be interpreted as the local ring of the adic point of $\mathbb{P}_K^{1,\text{an}}$ specializing x_1 in the direction towards 0.)

Let \mathcal{R}^{bd} be the subring of \mathcal{R} consisting of formal sums with bounded coefficients; these are exactly the elements of \mathcal{R} which define *bounded* analytic functions on $A(\rho, 1)$ for some $\rho \in (0, 1)$. The ring \mathcal{R}^{bd} carries a multiplicative *Gauss norm* defined by

$$\left| \sum_{i \in \mathbb{Z}} a_i z^i \right| = \max_i \{|a_i|\};$$

let \mathcal{R}^{int} be the subring of \mathcal{R}^{bd} consisting of elements of Gauss norm at most 1.

Lemma 10.3. *The ring \mathcal{R}^{int} is a Henselian discrete valuation ring. Consequently, \mathcal{R}^{bd} is a Henselian local field with residue field F .*

Proof. See [112, Proposition 3.2]. □

We next prepare to formulate the comparison between Kummer theory and Artin–Schreier–Witt theory by introducing the two sides of the comparison.

Definition 10.4. For any field L of characteristic not equal to p , for L^{sep} a separable closure of L , taking Galois cohomology on the exact sequence

$$1 \rightarrow \mu_{p^n} \rightarrow (L^{\text{sep}})^{\times} \xrightarrow{\bullet^{p^n}} (L^{\text{sep}})^{\times} \rightarrow 1$$

of G_L -modules gives the *Kummer isomorphism*

$$L^{\times} / L^{\times p^n} \cong H^1(G_L, \mu_{p^n})$$

because $H^1(G_L, (L^{\text{sep}})^{\times}) = 0$ by Noether’s form of Hilbert’s Theorem 90.

In the case $L = \mathcal{R}^{\text{bd}}$, by Lemma 10.3 we have a surjection $G_L \rightarrow G_F$ identifying G_F with the quotient of the maximal unramified extension of L . We thus obtain a restriction map $H^1(G_F, \mu_{p^n}) \rightarrow H^1(G_L, \mu_{p^n})$ and thus a map

$H^1(G_F, \mu_{p^n}) \rightarrow L^\times/L^{\times p^n}$. Note that G_F acts trivially on μ_{p^n} , so we may identify μ_{p^n} as a G_F -module with $\mathbb{Z}/p^n\mathbb{Z}$ by identifying our chosen primitive p^n -th root of unity $\zeta_{p^n} \in \mu_{p^n}$ with $1 \in \mathbb{Z}/p^n\mathbb{Z}$. We thus end up with a homomorphism

$$H^1(G_F, \mathbb{Z}/p^n\mathbb{Z}) \rightarrow (\mathcal{R}^{\text{bd}})^\times / (\mathcal{R}^{\text{bd}})^{\times p^n}. \quad (20)$$

(Beware that the opposite sign is used to normalize the isomorphism $\mu_{p^n} \cong \mathbb{Z}/p^n\mathbb{Z}$ in [113, 139], leading to some minor differences in the formulas.)

Definition 10.5. Consider the exact sequence

$$0 \rightarrow \mathbb{Z}/p^n\mathbb{Z} = W_n(\mathbb{F}_p) \rightarrow W_n(F^{\text{sep}}) \xrightarrow{1-\varphi} W_n(F^{\text{sep}}) \rightarrow 0$$

where φ denotes the Frobenius endomorphism of $W_n(F^{\text{sep}})$. The additive group $W_n(F^{\text{sep}})$ is a successive extension of copies of the additive group of F^{sep} ; since $H^1(G_F, F^{\text{sep}}) = 0$ by the additive version of Theorem 90, we also have $H^1(G_F, W_n(F^{\text{sep}})) = 0$. We thus obtain the *Artin–Schreier–Witt isomorphism*

$$\text{coker}(1 - \varphi, W_n(F)) \cong H^1(G_F, \mathbb{Z}/p^n\mathbb{Z}). \quad (21)$$

Combining this isomorphism with the map (20) derived from the Kummer isomorphism, we obtain a homomorphism

$$\text{coker}(1 - \varphi, W_n(F)) \rightarrow (\mathcal{R}^{\text{bd}})^\times / (\mathcal{R}^{\text{bd}})^{\times p^n}. \quad (22)$$

We note in passing how Swan conductors appear in the Artin–Schreier–Witt isomorphism. See also [81, Theorem 3.2], [70, Theorem 1.1], [158, Proposition 4.2]; see especially [28, Theorem 4.9] for the full computation of $\text{coker}(1 - \varphi, W_n(F))$.

Lemma 10.6. *For $\bar{a} \in W_n(F)$, let $\pi : G_F \rightarrow K^\times$ be the character corresponding via (21) to the class of \bar{a} in $\text{coker}(\varphi - 1, W_n(F))$. For $j = 0, \dots, n-1$, let m_j be the negation of the \bar{z} -adic valuation of \bar{a}_j , and assume that m_j is not a positive multiple of p . Then the Swan conductor of π equals*

$$\max\{0, m_0 p^{n-1}, m_1 p^{n-2}, \dots, m_{n-1}\}.$$

Proof. By hypothesis, if m_j is positive, then it is not divisible by p , so $m_j p^{n-1-i}$ has p -adic valuation $n - i - 1$. Consequently, any two of the quantities $m_0 p^{n-1}, m_1 p^{n-2}, \dots, m_{n-1}$, if they are nonzero, must be distinct. It thus suffices to check the claim in case \bar{a} is a Teichmüller element $[\bar{a}]$ for some $\bar{a} \in F$ of \bar{z} -adic valuation $-m$ for some integer m which is positive and not divisible by p . By splitting \bar{a} into powers of \bar{z} , we may further reduce to the case $\bar{a} = c\bar{z}^{-m}$. Using the compatibility of Swan conductors with tame base extensions, we may further reduce to the case $\bar{a} = \bar{z}^{-1}$.

For $j = 1, \dots, n$, let F_j be the extension of F obtained by adjoining the coordinates $\bar{b}_0, \dots, \bar{b}_{j-1}$ of a Witt vector \bar{b} satisfying

$$\bar{b} - \varphi(\bar{b}) = [\bar{a}].$$

It is clear that $\bar{b}_0 - \bar{b}_0^p = \bar{a}$. By direct computation, one sees that for $j = 1, \dots, n-1$, $\bar{b}_j - \bar{b}_j^p$ has the same \bar{z} -adic valuation as $\bar{b}_{j-1} \bar{a}^{p^j - p^{j-1}}$ and that this valuation is $-(p^j - p^{j-1} + \dots + p^{-j})$. It follows that the breaks in the lower numbering filtration of $\text{Gal}(F_n/F)$ occur at $(p^{2j+1} + 1)/(p + 1)$ for $j = 0, \dots, n-1$; by Herbrand's formula [154, Chap. IV], the breaks in the upper numbering filtration occur at $1, p, \dots, p^{n-1}$. The last of these breaks is the Swan conductor of π , proving the claim. \square

Corollary 10.7. *With notation as in Lemma 10.6, if b_j is the Swan conductor of $\pi^{\otimes p^{n-j}}$, then $b_j \geq pb_{j-1}$ for $j = 1, \dots, n$.*

Proof. The representation $\pi^{\otimes p^{n-j}}$ corresponds via (21) to the class of $p^{n-j}\bar{a}$ in $\text{coker}(\varphi - 1, W_j(F))$. By (11), this class is also represented by $V^{n-j}(\bar{a})$. We may now apply Lemma 10.6 to deduce the claim. \square

We now make explicit the relationship between the Kummer and Artin–Schreier–Witt isomorphisms.

Theorem 10.8 (Matsuda). *The homomorphism (22) is induced by a homomorphism*

$$W_n(\mathcal{R}^{\text{int}}) \rightarrow (\mathcal{R}^{\text{int}})^\times, \quad \underline{a} \mapsto E_{n,p}(p^n \underline{a}). \quad (23)$$

Proof. We first clarify the interpretation of the expression $E_{n,p}(p^n \underline{a})$ as an element of \mathcal{R}^{int} . Since by definition $E_{n,p}(\underline{a}) = \prod_{i=0}^{n-1} E_{n-i,p}(a_i)$, this amounts to evaluating the formal power series $E_{n,p}(p^n t)$ at $t = g$ for some $g = \sum_{i \in \mathbb{Z}} g_i z^i \in \mathcal{R}^{\text{int}}$. By Lemma 9.3, there exists $\rho_1 > 1$ such that the series $E_{n,p}(p^n t)$ converges for $|t| < \rho_1$. By the definition of \mathcal{R} , there exists $\rho_2 \in (0, 1)$ such that for any z in any nonarchimedean field containing K with $\rho_2 < |z| < 1$, $|g_i z^i| < \rho_1$ for all but finitely many $i < 0$. The same remains true if we replace ρ_2 by any larger value in $(0, 1)$; by so doing, we can force the inequality $|g_i z^i| < \rho_1$ to hold for all $i < 0$. The same inequality also holds for $i \geq 0$ because $|g_i| \leq 1$ by the definition of \mathcal{R}^{int} ; consequently, $g(z)$ belongs to the region of convergence of $E_{n,p}(p^n t)$, and the evaluation $E_{n,p}(p^n g)$ is well-defined as an analytic function on the annulus $A(\rho_2, 1)$. Since both $E_{n,p}(p^n t)$ and g have coefficients in \mathfrak{o}_K , the same is true of the composition, so $E_{n,p}(p^n g) \in \mathcal{R}^{\text{int}}$.

By (14), the map (23) is a homomorphism. Given \underline{a} , choose a minimal finite separable extension S of F such that there exists $\bar{b} \in W_n(S)$ with

$$\bar{b} - \varphi(\bar{b}) = \bar{a}.$$

(This amounts to forming a tower of Artin–Schreier extensions over F .) Apply Lemma 10.3 to construct a finite étale algebra \mathcal{S}^{int} over \mathcal{R}^{int} with residue field S . Choose a lift $\underline{b} \in \mathcal{S}^{\text{int}}$ of \bar{b} ; then $\underline{a} + \sigma(\underline{b}) - \underline{b} \in W_n(p\mathcal{S}^{\text{int}})$. By Lemma 9.10(b), we may define an element

$$f := F_{n,p}(\underline{b})E_{n,p}(\underline{a} + \sigma(\underline{b}) - \underline{b}) \in \mathcal{S}^{\text{int}}.$$

Then $f^{p^n} = E_{n,p}(p^n \underline{a})$.

On one hand, the image of \underline{a} in $\text{coker}(\varphi - 1, W_n(F))$ corresponds to the element of $H^1(G_F, \mathbb{Z}/p^n\mathbb{Z})$ which factors through $H^1(\text{Gal}(S/F), \mathbb{Z}/p^n\mathbb{Z})$ and sends $g \in \text{Gal}(S/F)$ to the integer $m \in \mathbb{Z}/p^n\mathbb{Z}$ for which $\varphi(\bar{b}) = \bar{b} + m$. On the other hand, we have $g(\underline{b}) = \underline{b} + m + \underline{c}$ for some $\underline{c} \in W((\zeta_{p^n} - 1)\mathcal{R}_n^{\text{int}})$, and so

$$\begin{aligned} g(f) &= F_{n,p}(\underline{b} + \underline{c} + m)E_{n,p}(\underline{a} + \sigma(\underline{b} + \underline{c} + m) - \underline{b} - \underline{c} - m) \\ &= F_{n,p}(\underline{b} + \underline{c} + m)F_{n,p}(\underline{b})^{-1} \\ &\quad E_{n,p}(\sigma(\underline{b} + \underline{c} + m) - \sigma(\underline{b}) - \underline{c} - m) \\ &\quad F_{n,p}(\underline{b})E_{n,p}(\underline{a} + \sigma(\underline{b}) - \underline{b}) \\ &= \zeta_{p^n}^m f \end{aligned}$$

by Lemma 9.11. It follows that (23) induces (22) as desired. \square

Remark 10.9. By comparing a character with its p -th power, we may deduce from Theorem 10.8 that

$$\frac{E_{n,p}(p^n \underline{a})}{E_{n-1,p}(p^{n-1} \underline{a})} \in (\mathcal{R}^{\text{int}})^{\times p^{n-1}}.$$

This may also be seen directly from Lemma 9.10 by rewriting the left side as $G_{n-1,p}(\underline{a})^{p^{n-1}}$.

11 Automorphisms of a Formal Disc

To conclude, we use Kummer–Artin–Schreier–Witt theory to translate the Oort local lifting problem into a question about the construction of suitable connections on \mathbb{P}_K^1 , and use this interpretation to describe existing combinatorial invariants connected with the Oort problem in terms of convergence polygons.

Hypothesis 11.1. Throughout Sect. 11, retain Hypothesis 10.1 and additionally fix $\underline{a} \in W_n(\mathcal{R}^{\text{int}})$.

Definition 11.2. Let $\pi : G_F \rightarrow \mu_{p^n}$ be the character corresponding to \underline{a} via the maps $W_n(\mathcal{R}^{\text{int}}) \rightarrow \text{coker}(\varphi - 1, W_n(F)) \cong H^1(G_F, \mu_{p^n})$ (the latter isomorphism being (21)). For $i = 1, \dots, n$, let b_i be the Swan conductor of $\pi^{\otimes p^{n-i}}$.

As before, let $x_1 \in \mathbb{P}_K^{1, \text{an}}$ denote the generic point of the disc $|z| < 1$. The residue field $\mathcal{H}(x_1)$ is the fraction field of a Cohen ring for the field $k(\bar{z})$. We have an embedding of $\mathcal{H}(x_1)$ into the completion of \mathcal{R}^{bd} for the Gauss norm, arising from an inclusion of Cohen rings lifting the inclusion $k(\bar{z}) \subset F$.

Apply the Katz–Gabber construction [83] to lift π to a representation of $G_{k(\bar{z})}$ unramified away from $\{0, \infty\}$, then identify the latter with a representation $\tilde{\pi} : G_{\mathcal{H}(x_1)} \rightarrow \mu_{p^n}$. By Crew’s analogue of the Katz–Gabber construction for p -adic differential equations [52], the character $\tilde{\pi}$ arises from a finite Galois cover of $A(\rho_1, \rho_2)$ for some $\rho_1 < 1 < \rho_2$. We may then proceed as in Definition 8.2 to obtain a rank 1 bundle \mathcal{E}_n with connection on this subspace. As in [139, Theorem 3.1] (modulo a sign convention; see Definition 10.4), this connection can be described explicitly as the free vector bundle on a single generator \mathbf{v} equipped with the connection

$$\nabla(\mathbf{v}) = \sum_{i=0}^{n-1} \sum_{j=0}^{n-i-1} (\zeta_{p^{n-j}} - 1) a_i^{p^j - 1} \mathbf{v} \otimes da_i.$$

Formally, we have $d\mathbf{v} = \mathbf{v} \otimes d \log E_{n,p}(\underline{a})$.

For $i = 1, \dots, n$, put $\mathcal{E}_i = \mathcal{E}_n^{\otimes p^{n-i}}$. Then \mathcal{E}_i corresponds to the character $\pi^{\otimes p^{n-i}}$ of order p^i in a similar fashion.

Remark 11.3. The construction given in Definition 11.2 is precisely the (φ, ∇) -module associated with π by the work of Fontaine and Tsuzuki [162]. In particular, it is unique in the sense that given any other construction, the two become isomorphic on $A(\rho_1, \rho_2)$ for some convenient choice of $\rho_1 < 1 < \rho_2$. Similarly, the restriction of the connection to $A(1, \rho_2)$ is the (φ, ∇) -module associated with the restriction of π to the decomposition group at ∞ , and enjoys a similar uniqueness property.

In fact, we can say something stronger. Recall that the Katz extension of a representation of G_F has only tame ramification at ∞ . Since π factors through a p -group, its Katz extension must in fact be unramified at ∞ , so it descends to a representation of the étale fundamental group of $\mathbb{P}_K^1 - \{0\}$. Such a representation gives rise to an overconvergent F -isocrystal, as shown by Crew [50, Theorem 3.1]; this means that for some choice of $\rho_1 < 1 < \rho_2$, the connection on $A(\rho_1, \rho_2)$ described above extends to the subspace $|z| > \rho_1$ of \mathbb{P}_K^1 .

Definition 11.4. Let S be the fixed field of $\ker(\pi)$; we may identify S with $k(\bar{u})$ for some parameter \bar{u} . The action of G_F defines a continuous k -linear action τ of μ_{p^n} on $k[\bar{u}]$. A *solution of the lifting problem* for π is a lifting of τ to a continuous \mathfrak{o}_K -linear action $\tilde{\tau}$ of μ_{p^n} on $\mathfrak{o}_K[\bar{u}]$.

Conjecture 11.5 (Oort). *A solution of the lifting problem exists for every π .*

A spectacular breakthrough on this problem has been made recently in work of Obus–Wewers and Pop [123, 138].

Theorem 11.6. *For fixed π , a solution of the lifting problem exists over some finite extension of K (that is, the lifting problem is solved if we do not insist on the field of definition).*

We will not say anything more here about the techniques used to prove Theorem 11.6. Instead, we describe an equivalence between solutions of the lifting problem for π and extensions of the connection on \mathcal{E}_n .

Definition 11.7. Suppose that $\tilde{\tau}$ is a solution of the lifting problem. Then $\tilde{\tau}$ gives rise to a finite Galois cover of the disc $|z| < 1$, which thanks to Remark 11.3 may be glued together with the cover from Definition 11.2 to give a finite ramified cover $f_n : Y_n \rightarrow X$ with $X = \mathbb{P}_K^1$, such that x_1 has a unique preimage in Y_n . (This cover is constructed *a priori* at the level of analytic spaces, but descends to a cover of schemes by rigid GAGA; see Remark 7.8.) For $i = 1, \dots, n$, let $f_i : Y_i \rightarrow X$ be the cover corresponding to $\rho^{p^{n-i}}$ in similar fashion, and let Z_i be the ramification locus of f_i ; also put $Z_0 = \emptyset$. For each $x \in Z_i - Z_{i-1}$, \mathcal{E}_n is regular at x with exponent m/p^{n-i+1} for some $m \in \mathbb{Z} - p\mathbb{Z}$.

Using the Riemann–Hurwitz formulas in characteristic 0 and p , we obtain the following relationship between the ramification of τ and of f_n .

Lemma 11.8. *With notation as in Definition 11.7, for $i = 1, \dots, n$,*

$$2 - 2g(Y_i) = 2p^i - \sum_{j=1}^i (p^j - p^{j-1})(\text{length}(Z_j) - \text{length}(Z_{j-1})) \quad (24)$$

$$= 2p^i - \sum_{j=1}^i (p^j - p^{j-1})b_j. \quad (25)$$

Consequently,

$$\text{length}(Z_i) = b_i + 1 \quad (i = 1, \dots, n). \quad (26)$$

Proof. The Riemann–Hurwitz formula for f_i asserts that

$$2 - 2g(Y_i) = \deg(f_i)(2 - 2g(\mathbb{P}_K^1)) - \sum_{z \in X} (\deg(f_i) - \text{length}(f_i^{-1}(z))).$$

For $z \in X$, we have $\deg(f_i) - \text{length}(f_i^{-1}(z)) = 0$ unless $z \in Z_i$. If $z \in Z_j - Z_{j-1}$ for some $j \in \{1, \dots, i\}$, then each preimage of z in Y_i is fixed by a group of order p^{i-j+1} , so $\text{length}(f_i^{-1}(z)) = p^{i-1}$. This yields (24).

Let $\bar{f}_i : \bar{Y}_i \rightarrow \mathbb{P}_k^1$ be the reduction of f_i ; then, $g(\bar{Y}_i) = g(Y_i)$. Set notation as in the proof of Lemma 10.6. Since \bar{f}_i is Galois and only ramifies above 0, the Riemann–Hurwitz formula for \bar{f}_i (see [156, Corollary 3.4.14, Theorem 3.8.7]) can be written in the form

$$2 - 2g(\bar{Y}_i) = \deg(f_i)(2 - 2g(\mathbb{P}_k^1)) - \sum_{m=0}^{\infty} (\#\text{Gal}(F_i/F)_m - 1)$$

where $\text{Gal}(F_i/F)_m$ denotes the m -th subgroup of $\text{Gal}(F_i/F)$ in the lower numbering filtration. By identifying these subgroups as in the proof of Lemma 10.6, we obtain (25). By combining (24) and (25), we may solve for b_i for $i = 1, \dots, n$ in succession to obtain (26). \square

We have the following explicit version of Theorem 5.5 in this setting.

Theorem 11.9. *Retain notation as in Definition 11.7. Let $\Gamma = \Gamma_{X, Z_n \cup \{\infty\}}$ be the union of the paths from ∞ to the elements of Z_n . Let \mathcal{N}_n be the convergence polygon of \mathcal{E}_n .*

- (a) *The function \mathcal{N}_n factors through the retraction of X^{an} onto Γ , and is affine on each edge of Γ .*
- (b) *Let $\Gamma_\infty \subset \Gamma$ be the path from x_1 to ∞ . For each $x \in \Gamma_\infty$, $s_1(\mathcal{N}_n(x)) = 0$.*
- (c) *The measure $\chi(\mathcal{E})$ (more precisely, the measure $\chi(\mathcal{E}, Z_n)$ in the notation of Remark 7.15) is discrete, supported at x_1 , of total measure $1 - b_n$.*
- (d) *For each $x \in \Gamma - \Gamma_\infty$, for \bar{i} the branch of x towards x_1 , $\partial_{\bar{i}} s_1(\mathcal{N}_n) = \ell - 1$ where ℓ is the length of the subset of Z_n dominated by x .*
- (e) *For $i = 1, \dots, n$, for each $x \in Z_i - Z_{i-1}$, let Γ_x be the pendant edge of Γ terminating at x . For each $y \in \Gamma_x$,*

$$s_1(\mathcal{N}_n(y)) = \left(n - i + 1 + \frac{1}{p-1} \right) \log p. \tag{27}$$

Proof. By Theorem 8.11(a), we have $s_1(\mathcal{N}_n(x_1)) = 0$; by Theorem 5.9(a), \mathcal{N}_n is identically zero on all x outside of the closed unit disc. In particular, we deduce (b).

By Theorem 7.23(c) and Remark 8.16, we must have $\chi_x(\mathcal{E}_n) \leq 0$ for all $x \in X^{\text{an}}$. By Theorem 7.7 (as reformulated in Theorem 7.12), Lemma 7.6, and Lemma 11.8, we have

$$\int_{X^{\text{an}}} \chi(\mathcal{E}_n) = \chi_{\text{dR}}(X^{\text{an}}, \mathcal{E}_n) = \chi_{\text{dR}}(U^{\text{an}}) = 2 - \text{length}(Z_n) = 1 - b_n.$$

By the previous paragraph plus Theorem 8.11(b), $\chi_{x_1}(\mathcal{E}) = 1 - b_n$; we thus deduce (c), which in turn immediately implies (a). Using Theorem 7.12 (which again applies in this situation thanks to Remark 8.16), we deduce (d). To obtain (e), in light of (a) we need only check (27) for y in some neighborhood of x . This follows by noting that as in Example 4.3, the binomial series $(1+z)^{1/p^n}$ has radius of convergence $p^{-n-1/(p-1)}$. \square

Remark 11.10. Retain notation as in Definition 11.7 and Theorem 11.9. The union of the paths from x_1 to the points of Z then forms a tree on which $s_1(\mathcal{N}_n)$ restricts to a harmonic function which is affine on each edge of the tree, with prescribed values and slopes at the pendant vertices (i.e., x_1 and the points of Z). However, the

existence of such a function imposes strong combinatorial constraints on the relative positions of the points of Z ; the resulting data are well-known in the literature on the Oort lifting problem, under the rubric of *Hurwitz trees* [24]. Similar data arise in [47, 157].

Remark 11.11. Conversely, suppose $X = \mathbb{P}_K^1$; Z equals $\{\infty\}$ plus a subset of the open unit disc; \mathcal{E} is a rank 1 vector bundle with connection on $U = X - Z$; for each $z \in Z$, ∇ is regular at z with exponent in $p^{-n}\mathbb{Z}$; and for some $\rho \in (0, 1)$, the restriction of \mathcal{E} to the space $|z| > \rho$ is isomorphic to \mathcal{E}_n . The connection $\mathcal{E}^{\otimes p^n}$ has only removable singularities, so it can be shown to be trivial either by passing to \mathbb{C} as in the proof of Lemma 7.6 and invoking complex GAGA (using the Riemann–Hilbert correspondence and the fact that $\mathbb{P}_{\mathbb{C}}^{1,\text{an}}$ is simply connected) or by applying the Dwork transfer theorem (Theorem 5.9(a)) to the disc $|z| < \sigma$ for some $\sigma > \rho$.

From this, we may see that \mathcal{E} arises as \mathcal{E}_n for some data as in Definition 11.2, i.e., that \mathcal{E} arises from a solution of the lifting problem. For example, to construct the finite map $f : Y \rightarrow X$, we may work analytically over \mathbb{C} (again as in the proof of Lemma 7.6), using complex GAGA to descend to the category of schemes. In the complex-analytic situation, we may locally choose a generator \mathbf{v} of \mathcal{E} , then form Y from X by adjoining $s^{1/p}$ where $sv^{\otimes p}$ is a nonzero horizontal section of $\mathcal{E}^{\otimes p^n}$. By similar considerations, we see that f is Galois with group μ_{p^n} and that $\mathcal{E} \cong \mathcal{E}_n$.

Remark 11.12. For a given π , it should be possible to construct a moduli space of solutions of the lifting problem in the category of rigid analytic spaces over K . Theorem 11.6 would then imply that this space is nonempty. Given this fact, it may be possible to derive additional results on the lifting problem, e.g., to resolve the case of dihedral groups. For $p > 2$, this amounts to showing that if τ anticommutes with the involution $\bar{z} \mapsto -\bar{z}$, then the action of the involution $z \mapsto -z$ fixes some point of the moduli space.

Acknowledgements Thanks to Francesco Baldassarri, Andrew Obus, Andrea Pulita, Junecue Suh, Daniele Turchetti, the anonymous referee, and the participants of the Simons Symposium for detailed feedback. The author was supported by NSF grant DMS-1101343 and UC San Diego (Warschawski Professorship), and by a Simons Visiting Professorship grant during November 2014 (visiting Pulita).

Appendix 1: Convexity

In this appendix, we give some additional technical arguments needed for the proofs of Theorems 6.6 and 7.23(c), which do not appear elsewhere in the literature. These arguments are not written in the same expository style as the rest of the main text; for instance, they assume much more familiarity with the author’s book [92].

Definition A.1. For I a subinterval of $[0, +\infty)$, let R_I be the ring of rigid analytic functions on the space $|t| \in I$ within the affine t -line over K , as in [97, Definition 3.1.1].

Hypothesis A.2. Throughout Appendix 1, assume $p = 0$, and let (M, D) be a differential module of rank n over $R_{[0, \beta)}$. For I a subinterval of $[0, \beta)$, write M_I as shorthand for $M \otimes_{R_{[0, \beta)}} R_I$.

Definition A.3. Let \mathbb{D}_β be the Berkovich disc $|t| < \beta$ over K . For $x \in \mathbb{D}_\beta$, define the real numbers $s_i(M, x)$ for $i = 1, \dots, n$ as in [97, Definition 4.3.2]; note that they are invariant under extension of K [97, Lemma 4.3.3]. For $r > -\log \beta$, define the functions

$$g_i(M, r) = -\log s_i(M, x_{e^{-r}})$$

$$G_i(M, r) = g_1(M, r) + \dots + g_i(M, r).$$

We begin with a variant of Theorem 7.23.

Lemma A.4. *The right slope of $G_n(M, r)$ at $r = -\log \beta$ equals $-\dim_K H^1(M)$, provided that at least one of the two is finite.*

Proof. Apply [135, Theorem 3.5.2]. □

Lemma A.5. *For $\rho \in (0, \beta)$, let x_ρ be the generic point of the disc $|t| \leq \rho$. Then for $i = 1, \dots, n$, the function $\rho \mapsto s_1(M, x_\rho) \cdots s_i(M, x_\rho)$ is nonincreasing in ρ .*

Proof. The case $i = n$ is immediate from Lemma A.4. To deduce the case $i < n$, it suffices to work locally around some $\rho_0 = e^{-r_0}$, and to check only those values of i for which $g_i(M, r_0) > g_{i+1}(M, r_0)$. For $\lambda \in K$, let M_λ be the differential module obtained from M by adding λ to D ; we may also view M_λ as the tensor product of M with the rank one differential module N_λ on a single generator \mathbf{v} satisfying $D(\mathbf{v}) = \lambda \mathbf{v}$. From the tensor product description, for $j = 1, \dots, n$ we have

$$g_j(M_\lambda, r) \leq \max\{g_j(M, r), g_1(N_\lambda, r)\}, \quad g_j(M, r) \leq \max\{g_j(M_\lambda, r), g_1(N_\lambda^\vee, r)\}.$$
(28)

In addition, by Example 4.1, we have

$$g_1(N_\lambda, r) = g_1(N_\lambda^\vee, r) = \max \left\{ -\log |\beta|, \frac{1}{p-1} \log p + \log |\lambda| \right\}$$

independently of r .

Choose an open subinterval I of $(g_{i+1}(M, r_0), g_i(M, r_0))$. After replacing K with a finite extension (which does not affect $G_j(M, r)$), we may choose λ in such a way that $g_1(N_\lambda, r) = g_1(N_\lambda^\vee, r)$ is equal to a constant value contained in I . We then have $g_i(M, r) > g_1(N_\lambda, r) > g_{i+1}(M, r)$ for all r in some neighborhood of r_0 . For such r , for $j \leq i$ we apply (28) to see that $g_j(M_\lambda, r) = g_j(M, r)$.

For $j > i$, (28) implies $g_j(M_\lambda, r) \leq g_1(N_\lambda, r)$, but we claim that in fact equality must hold. To wit, assume to the contrary that the inequality is strict; then, the restriction of M to the disc $|t| < e^{-g_j(M_\lambda, r)}$ would have a submodule isomorphic to N_λ^\vee . This would mean that M has a local horizontal section whose exact radius of

convergence is equal to $e^{-g_1(N_\lambda, r)}$, which would imply that there exists some value of k for which $g_k(M, r) = g_1(N_\lambda, r)$. However, this contradicts our choice of the interval I .

To summarize, for r in a neighborhood of r_0 we have

$$g_j(M_\lambda, r) = \begin{cases} g_j(M, r) & (j = 1, \dots, i) \\ g_j(N_\lambda, r) & (j = i + 1, \dots, n). \end{cases}$$

We thus deduce the original claim by applying the case $i = n$ to M_λ . \square

We now deduce Theorem 6.6.

Theorem A.6. *For $x, y \in \mathbb{D}_\beta$ such that x is the generic point of a disc containing y , for $i = 1, \dots, n$, we have*

$$s_1(M, x) \cdots s_i(M, x) \leq s_1(M, y) \cdots s_i(M, y).$$

In particular, the conclusion of Theorem 6.6 holds.

Proof. This follows from Lemma A.5 thanks to the invariance of the s_i under base extension. \square

We next proceed towards Theorem 7.23(c).

Definition A.7. Define the convergence polygon \mathcal{N} of M as in Definition 5.3, using the same disc at every point.

Lemma A.8. *Assume $p = 0$. For $i = 1, \dots, n$, for $x \in \mathbb{D}_\beta$, we have $\Delta h_i(\mathcal{N})_x \geq 0$.*

Proof. By base extension, we may reduce to the case $x = x_1$; we may also assume that the norm on K is nontrivial. As in the proof of Lemma A.5, we may further reduce to the case $i = n$. We may further reduce to the case where $H^0(M_{[0, \delta)}) = 0$ for all $\delta \in (1, \beta)$.

Choose a nonempty set W of K -rational points of \mathbb{D}_β with the following properties.

- (a) For each $w \in W$, $U_{w,1}$ is a branch of \mathbb{P}_K^1 at x_1 , which we also denote by \vec{t}_w .
- (b) The branches \vec{t}_w for $w \in W$ are pairwise distinct.
- (c) Let \vec{t}_∞ be the branch of \mathbb{P}_K^1 at x in the direction of ∞ . Then for all branches \vec{t} of \mathbb{P}_K^1 at x_1 , we have $\partial_{\vec{t}}(\mathcal{N}) = 0$ unless $\vec{t} = \vec{t}_\infty$ or $\vec{t} = \vec{t}_w$ for some $w \in W$.

Let $R_{w,I}$ be the ring of rigid analytic functions on the space $|z - w| \in I$, and put $M_{w,I} = M \otimes_{R_{[0, \beta)}} R_{w,I}$. Then by Lemma A.4, to prove the desired result, it suffices to check that for any $\gamma \in (0, 1)$, $\delta \in (1, \beta)$ sufficiently close to 1,

$$\dim_K H^1(M_{[0, \delta)}) \geq \sum_{w \in W} \dim_K H^1(M_{w, [0, \gamma)}).$$

It would hence also suffice to prove surjectivity of the map

$$H^1(M_{[0,\delta]}) \rightarrow \bigoplus_{w \in W} H^1(M_{w,[0,\gamma]}). \tag{29}$$

Since $p = 0$, we may invoke Theorem 7.12 to see that $H^1(M_{w,[0,\gamma]})$ is a finite-dimensional K -vector space, and so is complete with respect to its natural topology as a K -vector space (i.e., the one induced by the supremum norm with respect to some basis). Since the map $M_{w,[0,\gamma]} \rightarrow H^1(M_{w,[0,\gamma]})$ is a K -linear surjection from a Fréchet space over K to a Banach space over K , the Banach open mapping theorem implies that this map is a quotient map of topological spaces. In particular, since the map

$$M_{[0,\delta]} \rightarrow \bigoplus_{w \in W} M_{w,[0,\gamma]}$$

has dense image, so then does (29), proving the claim. □

Theorem A.9. *The conclusion of Theorem 7.23(c) holds.*

Proof. By Theorem 7.23(b), we may assume $x \notin \Gamma_{X,Z}$; the claim thus reduces to Lemma A.8. □

Remark A.10. In the proof of Lemma A.8, to deduce the finite-dimensionality of $H^1(M_{w,[0,\gamma]})$, one may replace the invocation of Theorem 7.12 with the following argument. For $\delta_1 \in (0, \gamma)$, $\delta_2 \in (\delta_1, \gamma)$ sufficiently large, by [97, Lemma 3.7.6] we have

$$\dim_K H^0(M_{w,(\delta_1,\gamma)}) = \dim_K H^1(M_{w,(\delta_1,\gamma)}), \quad \dim_K H^0(M_{w,(\delta_1,\delta_2)}) = \dim_K H^1(M_{w,(\delta_1,\delta_2)}).$$

By this calculation plus Mayer–Vietoris, the map $H^1(M_{w,[0,\gamma]}) \rightarrow H^1(M_{w,[0,\delta_2]})$ has finite-dimensional kernel and cokernel. Since $R_{w,[0,\gamma]} \rightarrow R_{w,[0,\delta_2]}$ is a compact map of topological K -vector spaces, the same is true of $M_{w,[0,\gamma]} \rightarrow M_{w,[0,\delta_2]}$; we may thus apply the Schwartz–Cartan–Serre lemma [100, Satz 1.2] to conclude.

Remark A.11. At this point, it is natural to consider what happens when $p > 0$. If we also add suitable hypotheses on p -adic non-Liouville exponents, then all of the preceding statements remain true (as in Remark 8.16). Absent such hypotheses, we cannot rely on finite-dimensionality of cohomology groups, but it may nonetheless be possible to adapt the proof of Theorem A.9 to the case $p > 0$ by establishing a relative version of Lemma A.4. To be precise, in the notation of the proof of Lemma A.8, one might hope to prove that

$$H^1(M_{[0,\delta]}/\bigoplus_{w \in W} M_{w,[0,\gamma]}) = 0$$

and to relate this vanishing directly to the Laplacian, bypassing the potential failure of finite-dimensionality for the individual cohomology groups.

Appendix 2: Thematic Bibliography

As promised in the introduction, we include an expansive but unannotated list of references related to key topics in the paper. For those topics discussed in [92], the chapter endnotes therein may be consulted for additional context.

- Underlying structure of Berkovich spaces: [8, 9], [6, Chap. 1], [18, 19, 25, 54–56, 76, 78, 130, 131, 166]. See also the survey [170] in this volume.
- Potential theory for Berkovich curves: [7, chapter by Baker], [6, 159].
- Formal structure of singular connections: [66, Chap. 4], [71], [82, Sect. 11], [92, Chap. 7], [93, 94, 102, 103, 108, 109, 118, 119], [155, Chap. 3], [165].
- Convergence of solutions of p -adic differential equations: [10, 13, 15, 16, 33, 36, 39, 45, 46, 64], [92, Chaps. 9–11], [99, 106, 132, 133, 136, 137, 140, 174]. For a retrospective circa 2000, see also [35].
- Transfer principles and effective convergence bounds: [34, 38, 63, 65], [92, 98, Chaps. 9, 13, 18].
- Logarithmic growth of p -adic solutions: [3, 29, 30, 33, 59, 60], [92, Chap. 18], [110, 124, 126].
- Stokes phenomena for complex connections: [104, 105, 107, 120, 148, 149], [155, Chaps. 7–9], [167, 168].
- Index formulas for nonarchimedean differential equations: [37, 40–42, 44, 86, 87, 113, 135, 142, 143, 145, 146, 161, 174]. See also the thematic bibliography of [147], and [43] for a survey of [40–42, 44].
- Comparison between algebraic and complex-analytic cohomology of connections: [4, 5, 53, 74, 115].
- Comparison between algebraic and nonarchimedean analytic cohomology of connections: [2, 4, 11, 12, 26, 27], [7, chapter by Kedlaya], [101].
- Decomposition theorems for nonarchimedean connections: [33, 37, 40–42, 44, 63], [92, Chap. 12], [97, 134, 142, 144].
- Monodromy for p -adic connections (Crew’s conjecture): [1, 51, 84, 85, 89–91], [92, Chaps. 20, 21], [95, 97, 116, 162, 163].
- p -adic Liouville numbers and p -adic exponents: [37, 40–42, 44, 46, 62], [92, Chap. 13], [97].
- Ramification of maps of Berkovich curves: [47, 67, 68, 79], [92, Chap. 19], [157].
- Measures of ramification and p -adic differential equations: [14, 28, 86, 88], [92, Chap. 19], [96, 111–114, 125, 161, 171–173].
- Kummer theory in mixed characteristic (Kummer–Artin–Schreier–Witt theory): [72, 112, 117, 128, 139, 151, 152] (unpublished; see [117] instead), [153, 160, 169].
- Oort lifting problem: [20, 22, 24, 31, 32, 49, 69, 72, 73], [77, Sect. 9], [121–123, 127–129, 138, 150]. See also the lecture notes [23], the introductions to [123, 138], and the PhD thesis [164].

References

1. André, Y.: Filtrations de type Hasse-Arf et monodromie p -adique. *Invent. Math.* **148**, 285–317 (2002)
2. André, Y.: Comparison theorems between algebraic and analytic de Rham cohomology. *J. Théor. Nombres Bordeaux* **16**, 335–355 (2004)
3. André, Y.: Dwork’s conjecture on the logarithmic growth of solutions of p -adic differential equations. *Compos. Math.* **144**, 484–494 (2008)
4. André, Y., Baldassarri, F.: *De Rham Cohomology of Differential Modules on Algebraic Varieties*. Progress in Mathematics, vol. 189. Birkhäuser, Basel (2001)
5. Atiyah, M.F., Hodge, W.V.D.: Integrals of the second kind on an algebraic variety. *Ann. Math.* **62**, 56–91 (1955)
6. Baker, M., Rumely, R.: *Potential Theory and Dynamics on the Berkovich Projective Line*. AMS Surveys and Monographs, vol. 159. American Mathematical Society, Providence (2010)
7. Baker, M., et al.: *p -adic Geometry: Lectures from the 2007 Arizona Winter School*. American Mathematical Society, Providence (2008)
8. Bayer, M., Payne, S., Rabinoff, J.: Nonarchimedean geometry, tropicalization, and metrics on curves. *J. Algebr. Geom.* (2015). arXiv:1104.0320v3
9. Baker, M., Payne, S., Rabinoff, J.: On the structure of non-Archimedean analytic curves. In: *Tropical and Non-Archimedean Geometry*. Contemporary Mathematics, vol. 605, pp. 93–121. American Mathematical Society, Providence, RI (2013)
10. Baldassarri, F.: Differential modules and singular points of p -adic differential equations. *Adv. Math.* **44**, 155–179 (1982)
11. Baldassarri, F.: Comparaison entre la cohomologie algébrique et la cohomologie p -adique rigide à coefficients dans un module différentiel, I. Cas des courbes. *Inv. Math.* **87**, 83–99 (1987)
12. Baldassarri, F.: Comparaison entre la cohomologie algébrique et la cohomologie p -adique rigide à coefficients dans un module différentiel, II. Cas des singularités régulières à plusieurs variables. *Math. An.* **280**, 417–439 (1988)
13. Baldassarri, F.: Continuity of the radius of convergence of differential equations on p -adic analytic curves. *Invent. Math.* **182**, 513–584 (2010)
14. Baldassarri, F.: Radius of convergence of p -adic connections and the p -adic Rolle theorem. *Milan J. Math.* **81**, 397–419 (2013)
15. Baldassarri, F., Di Vizio, L.: Continuity of the radius of convergence of p -adic differential equations on Berkovich spaces (2007). arXiv:07092008v1
16. Baldassarri, F., Kedlaya, K.S.: Harmonic functions attached to meromorphic connections on non-Archimedean curves (in preparation)
17. Berkovich, V.: Étale cohomology for non-Archimedean analytic spaces. *Publ. Math. IHÉS* **78**, 5–161 (1993)
18. Berkovich, V.: Smooth p -adic analytic spaces are locally contractible. *Invent. Math.* **137**, 1–84 (1999)
19. Berkovich, V.: Smooth p -adic analytic spaces are locally contractible, II. In: *Geometric Aspects of Dwork Theory*, vol. I, pp. 293–370. de Gruyter, Berlin (2004)
20. Bertin, J.: Obstructions locales, au relèvement de revêtements galoisiens de courbes lisses. *C.R. Acad. Sci. Paris Sér. I Math.* **326**, 55–58 (1998)
21. Beukers, F.: On Dwork’s accessory parameter problem. *Math. Z.* **241**, 425–444 (2002)
22. Bouw, I.I., Wewers, S.: The local lifting problem for dihedral groups. *Duke Math. J.* **134**, 421–452 (2006)
23. Bouw, I., Wewers, S.: Group actions on curves and the lifting problem. Course notes for the 2012 Arizona Winter School. Available at <http://math.arizona.edu/~swc/aws/2012/>
24. Brewis, L.H., Wewers, S.: Artin characters, Hurwitz trees, and the lifting problem. *Math. Ann.* **345**, 711–730 (2009)

25. Chambert-Loir, A., Ducros, A.: Formes différentielles réelles et courants sur les espaces de Berkovich. Preprint available at <http://www.math.jussieu.fr/~ducros/> (2016), downloaded in March 2016
26. Chiarellotto, B.: Sur le théorème de comparaison entre cohomologies de de Rham algébrique et p -adique rigide. *Ann. Inst. Fourier (Grenoble)* **38**, 1–15 (1988)
27. Chiarellotto, B.: A comparison theorem in p -adic cohomology. *Ann. Mat. Pura Appl.* **153**, 115–131 (1988)
28. Chiarellotto, B., Pulita, A.: Arithmetic and differential Swan conductors of rank 1 representations with finite local monodromy. *Am. J. Math.* **131**, 1743–1794 (2009)
29. Chiarellotto, B., Tsuzuki, N.: Logarithmic growth and Frobenius filtrations for solutions of p -adic differential equations. *J. Inst. Math. Jussieu* **8**, 465–505 (2009)
30. Chiarellotto, B., Tsuzuki, N.: Log-growth filtration and Frobenius slope filtration of F -isocrystals at the generic and special points. *Doc. Math.* **16**, 33–69 (2011)
31. Chinburg, T., Guralnick, R., Harbater, D.: Oort groups and local lifting problems. *Compos. Math.* **144**, 849–866 (2008)
32. Chinburg, T., Guralnick, R., Harbater, D.: The local lifting problem for actions of finite groups on curves. *Ann. Scient. Éc. Norm. Sup.* **44**, 537–605 (2011)
33. Christol, G.: Modules Différentielles et Équations Différentielles p -adiques. Queen’s University, Kingston (1983)
34. Christol, G.: Un théorème de transfert pour les disques singuliers réguliers, Cohomologie p -adique. *Astérisque* **119–120**, 151–168 (1984)
35. Christol, G.: Thirty years later. In: *Geometric Aspects of Dwork Theory*, vol. I, pp. 419–436. de Gruyter, Berlin (2004)
36. Christol, G.: The radius of convergence function for first order differential equations. In: *Advances in Non-Archimedean Analysis. Contemporary Mathematics*, vol. 551, pp. 71–89. American Mathematical Society, Providence, RI (2011)
37. Christol, G.: Le Théorème de Turrittin p -adique (in preparation); version of 11 Jun 2011 downloaded from <http://www.math.jussieu.fr/~christol>
38. Christol, G., Dwork, B.: Effective p -adic bounds at regular singular points. *Duke Math. J.* **62**, 689–720 (1991)
39. Christol, G., Dwork, B.: Modules différentiels sur les couronnes. *Ann. Inst. Fourier (Grenoble)* **44**, 663–701 (1994)
40. Christol, G., Mebkhout, Z.: Sur le théorème de l’indice des équations différentielles. *Ann. Inst. Four.* **43**, 1545–1574 (1993)
41. Christol, G., Mebkhout, Z.: Sur le théorème de l’indice des équations différentielles, II, *Ann. Math.* **146**, 345–410 (1997)
42. Christol, G., Mebkhout, Z.: Sur le théorème de l’indice des équations différentielles, III. *Ann. Math.* **151**, 385–457 (2000)
43. Christol, G., Mebkhout, Z.: p -adic differential equations. In: *Algebra and Number Theory (Fez). Lecture Notes in Pure and Applied Mathematics*, vol. 208, pp. 105–116. Dekker, New York (2000)
44. Christol, G., Mebkhout, Z.: Sur le théorème de l’indice des équations différentielles, IV. *Invent. Math.* **143**, 629–672 (2001)
45. Christol, G., Remmal, S.: Irregular p -adic linear differential equations. In: *Algebra and Number Theory (Fez). Lecture Notes in Pure and Applied Mathematics*, vol. 208, pp. 195–206. Dekker, New York (2000)
46. Clark, D.: A note on the p -adic convergence of solutions of linear differential equations. *Proc. Am. Math. Soc.* **17**, 262–269 (1966)
47. Cohen, A., Temkin, M., Trushin, D.: Morphisms of Berkovich curves and the different function (2014). arXiv:1408.2949v2
48. Conrad, B.: Relative ampleness in rigid geometry. *Ann. l’Institut Four. (Grenoble)* **56**, 1049–1126 (2006)
49. Corry, S.: Galois covers of the open p -adic disc. *Manuscripta Math.* **131**, 43–61 (2010)

50. Crew, R.: F -isocrystals and p -adic representations. In: Algebraic Geometry—Bowdoin 1985, Part 2. Proceedings of Symposia in Pure Mathematics, vol. 46.2, 111–138. American Mathematical Society, Providence (1987)
51. Crew, R.: Finiteness theorems for the cohomology of an overconvergent isocrystal on a curve. *Ann. Scient. Éc. Norm. Sup.* **31**, 717–763 (1998)
52. Crew, R.: Canonical extensions, irregularities, and the Swan conductor. *Math. Ann.* **316**, 19–37 (2000)
53. Deligne, P.: Équations Différentielles à Points Singuliers Réguliers. Lecture Notes in Mathematics, vol. 163. Springer, Berlin (1970)
54. Ducros, A.: Les espaces de Berkovich sont modérés, d’après E. Hrushovski and F. Loeser, Séminaire Bourbaki exposé 1056. *Astérisque* **352**, 459–507 (2013)
55. Ducros, A.: About Hrushovski and Loeser’s work on the homotopy type of Berkovich spaces (2015). arXiv:1309.0340v2
56. Ducros, A.: La Structure des Courbes Analytiques (in preparation); version of 12 Feb 2014 downloaded from <http://www.math.jussieu.fr/~ducros/livre.html>
57. Dwork, B.: On the zeta function of a hypersurface, II. *Ann. Math.* **80**, 227–299 (1964)
58. Dwork, B.: p -adic cycles. *Publ. Math. IHÉS* **37**, 27–115 (1969)
59. Dwork, B.: On p -adic differential equations, II: the p -adic asymptotic behavior of solutions of linear differential equations with rational function coefficients. *Ann. Math.* **98**, 366–376 (1973)
60. Dwork, B.: On p -adic differential equations, III: on p -adically bounded solutions of ordinary linear differential equations with rational function coefficients. *Invent. Math.* **20**, 35–45 (1973)
61. Dwork, B.: Bessel functions as p -adic functions of the argument. *Duke Math. J.* **41**, 711–738 (1974)
62. Dwork, B.: On exponents of p -adic differential equations. *J. Reine Angew. Math.* **484**, 85–126 (1997)
63. Dwork, B., Robba, P.: On ordinary linear p -adic differential equations, *Trans. Am. Math. Soc.* **231**, 1–46 (1977)
64. Dwork, B., Robba, P.: On natural radii of p -adic convergence. *Trans. Am. Math. Soc.* **256**, 199–213 (1979)
65. Dwork, B., Robba, P.: Effective p -adic bounds for solutions of homogeneous linear differential equations. *Trans. Am. Math. Soc.* **259**, 559–577 (1980)
66. Dwork, B., Gerotto, G., Sullivan, F.: An Introduction to G -Functions. *Annals of Mathematics Studies*, vol. 133. Princeton University Press, Princeton (1994)
67. Faber, X.: Topology and geometry of the Berkovich ramification locus for rational functions, I. *Manuscripta Math.* **142**, 439–474 (2013)
68. Faber, X.: Topology and geometry of the Berkovich ramification locus for rational functions, II. *Math. Ann.* **356**, 819–844 (2013)
69. Garuti, M.: Prolongement de revêtements galoisiens en géométrie rigide. *Compos. Math.* **104**, 305–331 (1996)
70. Garuti, M.: Linear systems attached to cyclic inertia. In: *Arithmetic Fundamental Groups and Noncommutative Algebra* (Berkeley, CA, 1999). Proceedings of Symposia in Pure Mathematics, vol. 70, 377–386. American Mathematical Society, Providence, RI (2002)
71. Gérard, R., Levelt, A.M.: Invariants mesurant l’irrégularité en un point singulier des systèmes d’équations différentielles linéaires. *Ann. Inst. Four. (Grenoble)* **23**, 157–195 (1973)
72. Green, B., Matignon, M.: Liftings of Galois covers of smooth curves. *Compos. Math.* **113**, 237–272 (1998)
73. Green, B., Matignon, M.: Order p automorphisms of the open disc of a p -adic field. *J. Am. Math. Soc.* **12**, 269–303 (1999)
74. Grothendieck, A.: On the de Rham cohomology of algebraic varieties. *Publ. Math. IHÉS* **29**, 95–103 (1966)
75. Grothendieck, A., et al.: SGA 1: Revêtements Étales et Groupe Fondamental, corrected edition, *Documents Mathématiques*, vol. 3. Société mathématique de France, Paris (2003)

76. Gubler, W., Rabinoff, J., Werner, A.: Skeletons and tropicalizations. *J. Adv. Math.* (2016). arXiv:1404.7044v3
77. Harbater, D., Obus, A., Pries, R., Stevenson, K.: Abhyankar's conjectures in Galois theory: current status and future directions (2014). arXiv:1408.0859v1
78. Hrushovski, E., Loeser, F.: Non-Archimedean tame topology and stably dominated types. In: *Annals of Mathematics Studies*. Princeton University Press, Princeton (2016)
79. Huber, R.: Swan representations associated with rigid analytic curves. *J. Reine Angew. Math.* **537**, 165–234 (2001)
80. Illusie, L.: Complexe de de Rham–Witt et cohomologie cristalline. *Ann. Sci. Éc. Norm. Sup.* **12**, 501–661 (1979)
81. Kato, K.: Swan conductors for characters of degree one in the imperfect residue field case. In: *Algebraic K -Theory and Algebraic Number Theory* (Honolulu, HI, 1987). *Contemporary Mathematics*, vol. 83, pp. 101–131. American Mathematical Society, Providence (1989)
82. Katz, N.M.: Nilpotent connections and the monodromy theorem: applications of a result of Turrittin. *Publ. Math. IHÉS* **39**, 175–232 (1970)
83. Katz, N.M.: Local-to-global extensions of representations of fundamental groups. *Ann. Inst. Four.* **36**, 69–106 (1986)
84. Kedlaya, K.S.: Semistable reduction for overconvergent F -isocrystals on a curve. *Math. Res. Lett.* **10**, 151–159 (2003)
85. Kedlaya, K.S.: A p -adic local monodromy theorem. *Ann. Math.* **160**, 93–184 (2004)
86. Kedlaya, K.S.: Local monodromy of p -adic differential equations: an overview. *Int. J. Number Theory* **1**, 109–154 (2005); errata posted at <http://kskedlaya.org/papers>
87. Kedlaya, K.S.: Fourier transforms and p -adic “Weil II”. *Compos. Math.* **142**, 1426–1450 (2006)
88. Kedlaya, K.S.: Swan conductors for p -adic differential modules, I: a local construction. *Alg. Number Theory* **1**, 269–300 (2007)
89. Kedlaya, K.S.: Semistable reduction for overconvergent F -isocrystals, I: unipotence and logarithmic extensions. *Compos. Math.* **143**, 1164–1212 (2007)
90. Kedlaya, K.S.: Semistable reduction for overconvergent F -isocrystals, II: a valuation-theoretic approach. *Compos. Math.* **144**, 657–672 (2008)
91. Kedlaya, K.S.: Semistable reduction for overconvergent F -isocrystals, III: local semistable reduction at monomial valuations. *Compos. Math.* **145**, 143–172 (2009)
92. Kedlaya, K.S.: p -Adic Differential Equations. Cambridge University Press, Cambridge (2010); errata posted at <http://kskedlaya.org/papers>
93. Kedlaya, K.S.: Good formal structures for flat meromorphic connections, I: surfaces. *Duke Math. J.* **154**, 343–418 (2010); erratum, *Duke Math. J.* **161**, 733–734 (2012)
94. Kedlaya, K.S.: Good formal structures for flat meromorphic connections, II: excellent schemes. *J. Am. Math. Soc.* **24**, 183–229 (2011)
95. Kedlaya, K.S.: Semistable reduction for overconvergent F -isocrystals, IV: local semistable reduction at nonmonomial valuations. *Compos. Math.* **147**, 467–523 (2011)
96. Kedlaya, K.S.: Swan conductors for p -adic differential modules, II: global variation. *J. Inst. Math. Jussieu* **10**, 191–224 (2011)
97. Kedlaya, K.S.: Local and global structure of connections on meromorphic curves. *Compos. Math.* **151**, 1096–1156 (2015)
98. Kedlaya, K.S., Tuitman, J.: Effective bounds for Frobenius structures. *Rend. Sem. Mat. Padova* **128**, 7–16 (2012)
99. Kedlaya, K.S., Xiao, L.: Differential modules on p -adic polyannuli. *J. Inst. Math. Jussieu* **9**, 155–201 (2010); erratum, *J. Inst. Math. Jussieu* **9**, 669–671 (2010)
100. Kiehl, R.: Der Endlichkeitsatz für eigentliche Abbildungen in der nichtarchimedischen Funktionentheorie. *Invent. Math.* **2**, 191–214 (1967)
101. Kiehl, R.: Die de Rham Kohomologie algebraischer Mannigfaltigkeiten über einem bewerteten Körper. *Publ. Math. IHÉS* **33**, 5–20 (1967)
102. Levelt, A.H.M.: Jordan decomposition for a class of singular differential operators. *Ark. Mat.* **13**, 1–27 (1975)

103. Levelt, A.H.M., van den Essen, A.R.P.: Irregular Singularities in Several Variables. *Memoirs of the American Mathematical Society*, vol. 40 (270). American Mathematical Society, Providence (1982)
104. Loday-Richaud, M.: Stokes phenomenon, multisummability and differential Galois groups. *Ann. Inst. Four. (Grenoble)* **44**, 849–906 (1994)
105. Loday-Richaud, M.: Solutions formelles des systèmes différentiels linéaires méromorphes et sommation. *Expo. Math.* **13**, 116–162 (1995)
106. Lutz, E.: Sur l'équation $y^2 = x^3 + Ax + B$ sur les corps p -adiques. *J. Reine Angew. Math.* **177**, 238–247 (1937)
107. Majima, H.: Asymptotic Analysis for Integrable Connections with Irregular Singular Points. *Lecture Notes in Mathematics*, vol. 1075. Springer, Berlin (1984)
108. Malgrange, B.: Sur les points singuliers des équations différentielles. *Enseign. Math.* **20**, 147–176 (1974)
109. Malgrange, B.: Connexions méromorphes 2: Le réseau canonique. *Invent. Math.* **124**, 367–387 (1996)
110. Manjra, S.: A note on non-Robba p -adic differential equations. *Proc. Jpn. Acad. Ser. A* **87**, 40–43 (2011)
111. Marmora, A.: Irrégularité et conducteur de Swan p -adiques. *Doc. Math.* **9**, 413–433 (2004)
112. Matsuda, S.: Local indices of p -adic differential operators corresponding to Artin-Schreier-Witt coverings. *Duke Math. J.* **77**, 607–625 (1995)
113. Matsuda, S.: Katz correspondence for quasi-unipotent overconvergent isocrystals. *Compos. Math.* **134**, 1–34 (2002)
114. Matsuda, S.: Conjecture on Abbes-Saito filtration and Christol-Mebkhout filtration. In: *Geometric Aspects of Dwork Theory*, vol. II, pp. 845–856. de Gruyter, Berlin (2004)
115. Mebkhout, Z.: Le théorème de comparaison entre cohomologies de de Rham d'une variété algébrique complexe et le théorème d'existence de Riemann. *Publ. Math. IHÉS* **69**, 47–89 (1989)
116. Mebkhout, Z.: Analogue p -adique du théorème de Turrittin et le théorème de la monodromie p -adique. *Invent. Math.* **148**, 319–351 (2002)
117. Mézard, A., Romagny, M., Tossici, D.: Sekiguchi-Suwa theory revisited. *J. Théor. Nombres Bordeaux* **26**, 163–200 (2014)
118. Mochizuki, T.: Good formal structure for meromorphic flat connections on smooth projective surfaces. In: *Algebraic Analysis and Around. Advanced Studies in Pure Mathematics*, vol. 54, pp. 223–253. Mathematical Society of Japan, Tokyo (2009)
119. Mochizuki, T.: Wild harmonic bundles and wild pure twistor D -modules. *Astérisque* **340** (2011)
120. Mochizuki, T.: The Stokes structure of a good meromorphic flat bundle. *J. Inst. Math. Jussieu* **10**, 675–712 (2011)
121. Obus, A.: The (local) lifting problem for curves. In: *Galois-Teichmüller Theory and Arithmetic Geometry. Advanced Studies in Pure Mathematics*, vol. 63, pp. 359–412. Mathematical Society of Japan, Tokyo (2012)
122. Obus, A.: A generalization of the Oort conjecture (2015). arXiv:1502.07623v1
123. Obus, A., Wewers, S.: Cyclic extensions and the local lifting problem. *Ann. Math.* **180**, 233–284 (2014)
124. Ohkubo, S.: A note on logarithmic growth Newton polygons of p -adic differential equations. *Int. Math. Res. Notices* (2014). Article ID rnu017
125. Ohkubo, S.: On differential modules associated to de Rham representations in the imperfect residue field case (2015). arXiv:1307.8110v4
126. Ohkubo, S.: On the rationality and continuity of logarithmic growth filtration of solutions of p -adic differential equations (2015). arXiv:1502.03804v2

127. Oort, F.: Lifting algebraic curves, abelian varieties, and their endomorphisms to characteristic zero. In: Algebraic Geometry, Bowdoin, 1985 (Brunswick, Maine, 1985). Proc. Symp. Pure Math., vol. 46, pp. 165–195. American Mathematical Society, Providence, RI (1987)
128. Oort, F., Sekiguchi, T., Suwa, N.: On the deformation of Artin-Schreier to Kummer. Ann. Scient. Éc. Norm. Sup. **22**, 345–375 (1989)
129. Pagot, G.: \mathbb{F}_p -espaces vectoriels de formes différentielles logarithmiques sur la droite projective. J. Number Theory **97**, 58–94 (2002)
130. Payne, S.: Analytification is the limit of all tropicalizations. Math. Res. Lett. **16**, 543–556 (2009)
131. Poineau, J.: Les espaces de Berkovich sont angéliques. Bull. Soc. Math. France **141**, 267–297 (2013)
132. Poineau, J., Pulita, A.: The convergence Newton polygon of a p -adic differential equation II: continuity and finiteness on Berkovich curves (2012). arXiv:1209.3663v1
133. Poineau, J., Pulita, A.: Continuity and finiteness of the radius of convergence of a p -adic differential equation via potential theory (2012). arXiv:1209.6276v1
134. Poineau, J., Pulita, A.: The convergence Newton polygon of a p -adic differential equation III: global decomposition and controlling graphs (2013). arXiv:1308.0859v1
135. Poineau, J., Pulita, A.: The convergence Newton polygon of a p -adic differential equation IV: local and global index theorems (2014). arXiv:1309.3940v2
136. Pons, E.: Polygone de convergence d'un module différentiel p -adique. C.R. Acad. Sci. Paris Sér. I Math. **327**, 77–80 (1998)
137. Pons, E.: Modules différentiels non solubles. Rayons de convergence et indices. Rend. Sem. Mat. Univ. Padova **103**, 21–45 (2000)
138. Pop, F.: Lifting of curves: the Oort conjecture. Ann. Math. **180**, 285–322 (2014)
139. Pulita, A.: Rank one solvable p -adic differential equations and finite Abelian characters via Lubin–Tate groups. Math. Ann. **337**, 489–555 (2007)
140. Pulita, A.: The convergence Newton polygon of a p -adic differential equation I: affinoid domains of the Berkovich affine line (2014). arXiv:1208.5850v4
141. Pulita, A.: Équations différentielles p -adiques, mémoire d'habilitation à diriger des recherches. Univ. Montpellier (2014)
142. Robba, P.: On the index of p -adic differential operators, I. Ann. Math. **101**, 280–316 (1975)
143. Robba, P.: On the index of p -adic differential operators, II. Duke Math. J. **43**, 19–31 (1976)
144. Robba, P.: Lemmes de Hensel pour les opérateurs différentiels. Application à la réduction formelle des équations différentielles. Ens. Math. **26**, 279–311 (1980)
145. Robba, P.: On the index of p -adic differential operators, III. Application to twisted exponential sums. Astérisque **119–120**, 191–266 (1984)
146. Robba, P.: Indice d'un opérateur différentiel p -adique, IV: Cas des systèmes. Mesure de l'irrégularité dans un disque. Ann. Inst. Fourier, Grenoble **35**, 13–55 (1985)
147. Robba, P.: Une introduction naïve aux cohomologies de Dwork. Mém. Soc. Math. France **23**, 61–105 (1986)
148. Sabbah, C.: Équations différentielles à points singuliers irréguliers et phénomène de Stokes en dimension 2. Astérisque **263** (2000)
149. Sabbah, C.: Introduction to Stokes Structures. Lecture Notes in Mathematics, vol. 2060. Springer, Berlin (2013)
150. Saïdi, M.: Fake liftings of Galois covers between smooth curves. In: Galois–Teichmüller Theory and Arithmetic Geometry. Advanced Studies in Pure Mathematics, vol. 63, pp. 457–501. Mathematical Society of Japan, Tokyo (2012)
151. Sekiguchi, T., Suwa, N.: Théorie de Kummer–Artin–Schreier et applications. J. Théorie Nombres Bordeaux **7**, 177–189 (1995)
152. Sekiguchi, T., Suwa, N.: On the unified Kummer–Artin–Schreier–Witt theory, Prépublications du laboratoire de Mathématiques Pures de Bordeaux, no. 111 (1999)
153. Sekiguchi, T., Suwa, N.: A note on extensions of algebraic and formal groups. IV. Kummer–Artin–Schreier–Witt theory of degree p^2 . Tohoku Math. J. **53**, 203–240 (2001)
154. Serre, J.-P.: Local Fields. Graduate Texts in Mathematics, vol. 67. Springer, New York (1979)

155. Singer, M., van der Put, M.: Galois Theory of Linear Differential Equations. Grundlehren der math. Wissenschaften, vol. 328. Springer, Berlin (2003)
156. Stichtenoth, H.: Algebraic Function Fields and Codes. Graduate Texts in Mathematics, vol. 254. Springer, Berlin (2009)
157. Temkin, M.: Metric uniformization of morphisms of Berkovich curves (2014). arXiv:1410.6892v1
158. Thomas, L.: Ramification groups in Artin-Schreier-Witt extensions. J. Théor. Nombres Bordeaux **17**, 689–720 (2005)
159. Thuillier, A.: Théorie du potentiel sur les courbes en géométrie analytique non archimédienne. PhD thesis, Université de Rennes (2005)
160. Tossici, D.: Models of $\mu_{p^2, K}$ over a discrete valuation ring (with an appendix by X. Caruso). J. Algebra **323**, 1908–1957 (2010)
161. Tsuzuki, N.: The local index and the Swan conductor. Comp. Math. **111**, 245–288 (1998)
162. Tsuzuki, N.: Finite local monodromy of overconvergent unit-root F -isocrystals on a curve. Am. J. Math. **120**, 1165–1190 (1998)
163. Tsuzuki, N.: Morphisms of F -isocrystals and the finite monodromy theorem for unit-root F -isocrystals. Duke Math. J. **111**, 385–419 (2002)
164. Turchetti, D.: Contributions to arithmetic geometry in mixed characteristic: lifting covers of curves, non-Archimedean geometry and the ℓ -modular Weil representation. PhD thesis, Univ. de Versailles St.-Quentin (2014). Available online at <http://webusers.imj-prg.fr/~daniele.turchetti/>
165. Turrittin, H.L.: Convergent solutions of ordinary linear homogeneous differential equations in the neighborhood of an irregular singular point. Acta Math. **93**, 27–66 (1955)
166. Tyomkin, I.: Tropical geometry and correspondence theorems via toric stacks. Math. Ann. **353**, 945–995 (2012)
167. Varadarajan, V.S.: Linear meromorphic differential equations: a modern point of view. Bull. Am. Math. Soc. **33**, 1–42 (1996)
168. Wasow, W.: Asymptotic Expansions for Ordinary Differential Equations, reprint of the 1976 edition. Dover, New York (1987)
169. Waterhouse, W.: A unified Kummer-Artin-Schreier sequence. Math. Ann. **277**, 447–451 (1987)
170. Werner, A.: Analytification and tropicalization over non-Archimedean fields. Simons Symposium 2015 lecture notes available at <http://users.math.yale.edu/~sp547/SimonsSymposium2015.html>
171. Xiao, L.: On ramification filtrations and p -adic differential equations, I: equal characteristic case. Alg. Number Theory **4**, 969–1027 (2010)
172. Xiao, L.: On ramification filtrations and p -adic differential equations, II: mixed characteristic case. Compos. Math. **148**, 415–463 (2012)
173. Xiao, L.: On the refined ramification filtrations in the equal characteristic case. Alg. Number Theory **6**, 1579–1667 (2012)
174. Young, P.T.: Radii of convergence and index for p -adic differential operators. Trans. Am. Math. Soc. **333**, 769–785 (1992)

About Hrushovski and Loeser’s Work on the Homotopy Type of Berkovich Spaces

Antoine Ducros

Abstract Using sophisticated model-theoretic tools, Hrushovski and Loeser have proved that the Berkovich analytification of a quasi-projective algebraic variety is locally contractible and has the homotopy type of a finite simplicial complex. This text is a survey of their work. We begin with a presentation of the model-theoretic notions involved (definability, types, etc.), and then give a sketch of their proof. At the end, we explain how one of their intermediate results can be applied to describe the pre-image of the skeleton of a torus under a finite map.

Keywords Berkovich analytic spaces • Skeletons • Definability • Model-theoretic types

1 Introduction

(1.1) At the end of the 1980s, Berkovich developed in [2, 3] a new theory of analytic geometry over non-Archimedean fields, coming after those by Krasner, Tate [22], and Raynaud [20]. One of the main advantages of this approach is that the resulting spaces enjoy very nice *topological* properties: they are locally compact and locally *pathwise* connected (although non-Archimedean fields are totally disconnected, and often not locally compact). Moreover, Berkovich spaces have turned out to be “tame” objects—in the informal sense of Grothendieck’s *Esquisse d’un programme*. Before illustrating this rather vague assertion by several examples, let us fix some terminology and notations.

(1.2) Let k be a field which is complete with respect to a non-Archimedean absolute value $|\cdot|$ and let k° be the valuation ring $\{x \in k, |x| \leq 1\}$.

(1.2.1) A (Berkovich) analytic space over k is a (locally compact, locally path-connected) topological space X equipped with some extra-data, among which:

A. Ducros (✉)
IMJ-PRG, F-75252 Paris Cedex 05, France
Institut Universitaire de France, Paris, France
e-mail: antoine.ducros@imj-prg.fr

- a sheaf of k -algebras whose sections are called *analytic functions*, which makes X a locally ringed space;
- for every point x of X , a complete non-Archimedean field $(\mathcal{H}(x), |\cdot|)$ endowed with an isometric embedding $k \hookrightarrow \mathcal{H}(x)$, and an “evaluation” morphism from $\mathcal{O}_{X,x}$ to $\mathcal{H}(x)$, which is denoted by $f \mapsto f(x)$.

(1.2.2) To every k -scheme of finite type X is associated in a functorial way a k -analytic space X^{an} , which is called its *analytification* [3, Sect. 2.6]; it is provided with a natural morphism of locally ringed spaces $X^{\text{an}} \rightarrow X$. If f is a section of \mathcal{O}_X , we will often still denote by f its pullback to $\mathcal{O}_{X^{\text{an}}}(X^{\text{an}})$.

Let us recall what the underlying topological space of X^{an} is. As a set, it consists of couples $(\xi, |\cdot|)$ where $\xi \in X$ and where $|\cdot|$ is a non-Archimedean absolute value on $\kappa(\xi)$ extending the given absolute value of k . If $x = (\xi, |\cdot|)$ is a point of X^{an} , then $\mathcal{H}(x)$ is the completion of the valued field $(\kappa(\xi), |\cdot|)$. The map $X^{\text{an}} \rightarrow X$ is nothing but $(\xi, |\cdot|) \mapsto \xi$; note that the pre-image of a Zariski-open subset U of X on X^{an} can be identified with U^{an} . The set X^{an} is equipped with the coarsest topology such that:

- $X^{\text{an}} \rightarrow X$ is continuous (otherwise said, U^{an} is an open subset of X^{an} for every Zariski-open subset U of X);
- for every Zariski-open subset U of X and every $f \in \mathcal{O}_X(U)$, the map $(\xi, |\cdot|) \mapsto |f(\xi)|$ from U^{an} to \mathbb{R}_+ is continuous.

Let us now assume that X is affine, say $X = \text{Spec } A$. The topological space underlying X^{an} can then be given another description: it is the set of all multiplicative maps $\varphi: A \rightarrow \mathbb{R}_+$ that extend the absolute value of k and that satisfy the inequality $\varphi(a + b) \leq \max(\varphi(a), \varphi(b))$ for every $(a, b) \in A^2$ (such a map will be simply called a *multiplicative semi-norm*); its topology is the one inherited from the product topology on \mathbb{R}_+^A . This is related as follows to the former description:

- to any couple $(\xi, |\cdot|)$ as above corresponds the multiplicative semi-norm φ defined by the formula $\varphi(f) = |f(\xi)|$;
- to any multiplicative semi-norm φ corresponds the couple $(\mathfrak{p}, |\cdot|)$ where \mathfrak{p} is the kernel of φ (this is a prime ideal of A) and $|\cdot|$ is the absolute value on $\text{Frac } A/\mathfrak{p}$ extending the multiplicative norm on A/\mathfrak{p} induced by φ .

Modulo this new description the map $X^{\text{an}} \rightarrow X$ simply sends a semi-norm to its kernel.

Definition 1.2.3. Let X be a k -scheme of finite type. A subset of X^{an} is said to be *semi-algebraic* if it can be defined, locally for the Zariski-topology of X , by a boolean combination of inequalities $|f| \bowtie \lambda |g|$ where f and g are sections of \mathcal{O}_X , where $\bowtie \in \{<, >, \leq, \geq\}$, and where $\lambda \in \mathbb{R}_+$; it is called *strictly semi-algebraic* if the λ 's can be chosen in $|k|$.

(1.2.4) To every k° -formal scheme \mathfrak{X} (topologically) of finite type is associated in a functorial way a compact k -analytic space \mathfrak{X}_η , which is called its *generic fiber* [4, Sect. 1].

In what follows, we will be interested in the generic fiber \mathfrak{X}_η for \mathfrak{X} being a *polystable* k° -formal scheme. This notion has been introduced by Berkovich

(Definition 1.2 of Berkovich [5]); let us simply emphasize here that semi-stable k^o -formal schemes are polystable.

Definition 1.2.5. A *polyhedron* is a subset of \mathbb{R}^n (for some n) which is a finite union of *rational* simplices.

(1.3) Examples of tameness properties of Berkovich spaces.

(1.3.1) Let \mathfrak{X} be a polystable k^o -formal scheme. Berkovich proves in [5] that its generic fiber \mathfrak{X}_η admits a strong deformation retraction to one of its closed subset $S(\mathfrak{X})$ which is homeomorphic to a polyhedron¹ of dimension $\leq \dim \mathfrak{X}$.

(1.3.2) Smooth k -analytic spaces are locally contractible; this is also proved in [5] by Berkovich, by reduction to Assertion (1.3.1) above through de Jong’s alterations.

(1.3.3) Let X be a k -scheme of finite type. Every semi-algebraic subset of X^{an} has finitely many connected components, each of which is semi-algebraic; this was proved by the author in [9].

(1.3.4) Let X be a compact k -analytic space and let f be an analytic function on X ; for every $\varepsilon \geq 0$, denote by X_ε the set of $x \in X$ such that $|f(x)| \geq \varepsilon$. There exists a finite partition \mathcal{P} of \mathbb{R}_+ in intervals such that for every $I \in \mathcal{P}$ and every $(\varepsilon, \varepsilon') \in I^2$ with $\varepsilon \leq \varepsilon'$, the natural map $\pi_0(X_{\varepsilon'}) \rightarrow \pi_0(X_\varepsilon)$ is bijective. This has been established by Poineau in [17] (it had already been proved in the particular case where f is invertible by Abbes and Saito in [1]).

(1.4) Hrushovski and Loeser’s work. As far as tameness is concerned, there has been in 2009–2010 a major breakthrough, namely, the work [16], by Hrushovski and Loeser. Let us quickly explain what their main results consist of. We fix a *quasi-projective* k -scheme and a semi-algebraic subset V of X^{an} . Hrushovski and Loeser have proven the following.

(1.4.1) There exists a strong deformation retraction from V to one of its closed subsets which is homeomorphic to a polyhedron. More precisely, Hrushovski and Loeser build a continuous map $h : [0; 1] \times V \rightarrow V$, and prove the existence of a compact subset S of V homeomorphic to a polyhedron such that the following hold:

- (i) $h(0, v) = v$ and $h(1, v) \in S$ for all $v \in V$;
- (ii) $h(t, v) = v$ for all $t \in [0; 1]$ and all $v \in S$;
- (iii) $h(1, h(t, v)) = h(1, v)$ for all $t \in [0; 1]$ and all $v \in V$.

(1.4.2) The topological space V is locally contractible.

(1.4.3) Let Y be a k -scheme of finite type and let $\varphi : X \rightarrow Y$ be a k -morphism. The set of homotopy types of fibers of the map $\varphi^{\text{an}}|_V : V \rightarrow Y^{\text{an}}$ is finite.²

¹It is more precisely homeomorphic to the (abstract) “incidence polyhedron” of the special fiber \mathfrak{X}_s that encodes the combinatorics of its singularities, e.g., it is reduced to a point if \mathfrak{X}_s is smooth and irreducible.

²One could ask the same question concerning the set of *homeomorphism* types; to the author’s knowledge, the answer is not known, and likely not easy to obtain using Hrushovski and Loeser’s approach.

(1.4.4) Let f be a function belonging to $\mathcal{O}_X(X)$. For every $\varepsilon \geq 0$ let us denote by V_ε the set of $x \in V$ such that $|f(x)| \geq \varepsilon$. There exists a finite partition \mathcal{P} of \mathbb{R}_+ in intervals such that for every $I \in \mathcal{P}$ and every $(\varepsilon, \varepsilon') \in I^2$ with $\varepsilon \leq \varepsilon'$, the embedding $V_{\varepsilon'} \hookrightarrow V_\varepsilon$ is a homotopy equivalence.

(1.5) Comments.

(1.5.1) Since V has a basis of open subsets which are semi-algebraic subsets of X^{an} , Assertion (1.4.2) is an obvious corollary of Assertion (1.4.1) and of the local contractibility of polyhedra; note that the property (iii) of (1.4.2) is here crucial.

(1.5.2) Even when X is smooth, projective and when $V = X^{\text{an}}$, Assertion (1.4.1) was previously known only when X has a polystable model.

(1.5.3) The quasi-projectivity assumption on X is needed by Hrushovski and Loeser for *technical* reasons, but everything remains likely true for X any k -scheme of finite type.

(1.5.4) In order to prove these tameness results, Hrushovski and Loeser develop a new kind of geometry over valued fields of arbitrary height, using highly sophisticated tools from model theory. In fact, most of their work is actually devoted to this geometry, and they transfer their results to the Berkovich framework only at the very end of their paper (Chap. 14 of Hrushovski and Loeser [16]); for instance, they first prove counterparts of Assertions (1.4.1)–(1.4.4) in their setting.

(1.6) About this text. In this article we survey the methods and results of [16].³ We will describe roughly the geometry of Hrushovski and Loeser, say a few words about the links between their spaces and those of Berkovich, and give a coarse sketch of the proof of their version of Assertion (1.4.1).

In the last section, we will explain how a key finiteness result of [16] has been used by the author in [10], to show the following: if X is an n -dimensional analytic space and if $f : X \rightarrow \mathbb{G}_{m,k}^{\text{an}}$ is any morphism, the pre-image of the “skeleton” of $\mathbb{G}_{m,k}^{\text{an}}$ under f inherits a canonical piecewise-linear structure.

2 Model Theory of Valued Fields: Basic Definitions

(2.1) General conventions about valued fields.

(2.1.1) In this text, a *valuation* will be an abstract Krull valuation, neither assumed to be of height one nor nontrivial. We will use the *multiplicative notation*: if k is a valued field, then unless otherwise stated its valuation will be denoted by $|\cdot|$, and its value group will not be given a specific name—we will simply use the notation $|k^\times|$, and 1 for its unit element. We will formally add to the ordered group $|k^\times|$ an

³The reader may also refer to the more detailed survey [12].

absorbing, smallest element 0 and set $|0| = 0$; with these conventions one then has $|k| = \{0\} \cup |k^\times|$ and

$$|a + b| \leq \max(|a|, |b|)$$

for every $(a, b) \in k^2$.

We will denote by $|k^\times|^\mathbb{Q}$ the divisible hull of the ordered group $|k^\times|$, and set for short $|k|^\mathbb{Q} = \{0\} \cup |k^\times|^\mathbb{Q}$.

The reader should be aware that the multiplicative notation, which is usual in Berkovich’s theory, is *not* consistent with Hrushovski and Loeser’s choice. Indeed, they use the additive notation, call “0” what we call “1”, call “ $+\infty$ ” what we call “0”, and their order is the opposite of ours.

(2.1.2) If k is a valued field, we will set

$$k^\circ = \{x \in k, |x| \leq 1\} \text{ and } k^{\circ\circ} = \{x \in k, |x| < 1\}.$$

Then k° is a subring of k ; it is local, and $k^{\circ\circ}$ is its unique maximal ideal. The quotient $k^\circ/k^{\circ\circ}$ will be denoted by \bar{k} ; it is called the *residue field* of $(k, |\cdot|)$ (or simply of k if there is no ambiguity about the valuation).

(2.2) Definable functors: the embedded case. Let k be a valued field. We fix an algebraic closure \bar{k} of k , and an extension of the valuation of k to \bar{k} . Let \mathbf{M} be the category of algebraically closed valued extensions of \bar{k} whose valuation is nontrivial (morphisms are isometric \bar{k} -embeddings).

(2.2.1) Definable subsets. Let X and Y be k -schemes of finite type and let $F \in \mathbf{M}$. A subset of $X(F)$ will be said to be *k-definable* if it can be defined, locally for the Zariski-topology of X , by a boolean combination of inequalities of the form

$$|f| \bowtie \lambda |g|$$

where f and g are regular functions, where $\lambda \in |k|$, and where $\bowtie \in \{\leq, \geq, <, >\}$. A map from a k -definable subset of $X(F)$ to a k -definable subset of $Y(F)$ will be said to be *k-definable* if its graph is a k -definable subset of $X(F) \times Y(F) = (X \times_k Y)(F)$.

(2.2.2) Comments. Definability is in fact a general notion of model theory which makes sense with respect to any given *language*. The one we have introduced here is actually definability in the *language of valued fields*, and the reader is perhaps more familiar with the two following examples: definability in the language of fields, which is nothing but Zariski-constructibility; and definability in the language of ordered fields, which is nothing but semi-algebraicity.

(2.2.3) Definable sub-functors. The scheme X induces a functor $F \mapsto X(F)$ from \mathbf{M} to **Sets**, which we will also denote by X . A sub-functor D of X will be said to be *k-definable* if it can be defined, locally for the Zariski-topology of X , by a boolean combination of inequalities of the form

$$|f| \bowtie \lambda |g|$$

where f and g are regular functions, where $\lambda \in |k|$, and where $\bowtie \in \{\leq, \geq, <, >\}$. A natural transformation from a k -definable sub-functor of X to a k -definable sub-functor of Y will be said to be k -definable if its graph is a k -definable sub-functor of $X \times_k Y$. If D is a k -definable sub-functor of X , then $D(F)$ is for every $F \in \mathbf{M}$ a k -definable subset of $X(F)$; if D' is a k -definable sub-functor of Y and if $f : D \rightarrow D'$ is a k -definable natural transformation, $f(F) : D(F) \rightarrow D'(F)$ is for every $F \in \mathbf{M}$ a k -definable map.

The following fundamental facts are straightforward consequences of the so-called *quantifier elimination for nontrivially valued algebraically closed fields*. It is usually attributed to Robinson, though the result proved in his book [21] is weaker—but the main ideas are there; for a complete proof, cf. [7, 19] or [23].

- (i) For every $F \in \mathbf{M}$, the assignment $D \mapsto D(F)$ establishes a one-to-one correspondence between k -definable *sub-functors* of X and k -definable *subsets* of $X(F)$. The analogous result holds for k -definable natural transformations and k -definable maps.
- (ii) If $f : D \rightarrow D'$ is a k -definable natural transformation between two k -definable sub-functors $D \subset X$ and $D' \subset Y$, the sub-functor of D' that sends F to the image of $f(F) : D(F) \rightarrow D'(F)$ is k -definable.

Comments about assertion (i). For a given $F \in \mathbf{M}$, assertion (i) allows to identify k -definable subsets of $X(F)$ with k -definable sub-functors of X , and the choice of one of those viewpoints can be somehow a matter of taste. But even if one is actually only interested in F -points, it can be useful to think of a k -definable subset D of $X(F)$ as a sub-functor of X . Indeed, this allows to consider points of D over valued fields larger than F , which often encode in a natural and efficient way the “limit” behavior of D .

(2.2.4) Subvarieties of X and definable sub-functors. Here we consider only *field-valued* points; for that reason, this theory is totally insensitive to nilpotents. For example, X_{red} and X define the same functor in our setting. But let us now give a more subtle example of this phenomenon. If $Y \hookrightarrow X$ is a k -subscheme, then $F \mapsto Y(F)$ is obviously a k -definable sub-functor of X .

Now, let L be the perfect closure of k (which embeds canonically in any $F \in \mathbf{M}$), and let Y be an L -subscheme of X_L . The functor $F \mapsto Y(F)$ is then again k -definable, even if Y cannot be defined over k in the usual, algebro-geometric sense. Indeed, to see it we immediately reduce to the case where X is affine; now, Y is defined by an ideal $(f_i)_{1 \leq i \leq n}$, where the f_i 's all belong to $\mathcal{O}_{X_L}(X_L)$. But then, again since we only consider field-valued points, the sub-functor $F \mapsto Y(F)$ is equal to the one defined by the ideal $(f_i^{p^n})$ for any n , where p is the characteristic exponent of k . By taking n large enough so that $f_i^{p^n} \in \mathcal{O}_X(X)$ for every i , we see that our functor is k -definable.

(2.3) Definable functors: the abstract case.

(2.3.1) Definably embeddable functors. Let us say that a functor D from \mathbf{M} to \mathbf{Sets} is *k -definably embeddable* if there exist a k -scheme of finite type X , a definable sub-functor D_0 of X , and an isomorphism $D \simeq D_0$.

The functor D_0 is not *a priori* uniquely determined (up to unique k -definable isomorphism). But for all functors D we will consider below, we will be given *implicitly* not only $D(F)$ for every $F \in \mathbf{M}$, but also $D(T)$ for every $F \in \mathbf{M}$ and every F -definable sub-functor T of an F -scheme of finite type; otherwise said, if D classifies objects of a certain kind, we not only know what such an object defined over F is, but also what an embedded F -definable family of such objects is. It follows from Yoneda’s lemma that such data ensure the canonicity of D_0 as soon as it exists (see [12, Sect. 1] for more detailed explanations about those issues).

(2.3.2) It turns out that the class of k -definably embeddable functors is too restrictive, for the following reason. Let D be such a functor and let R be a sub-functor of $D \times D$. Assume that R is itself k -definably embeddable, and that it is an equivalence relation—that is, that $R(F) \subset D(F) \times D(F)$ is the graph of an equivalence relation on $D(F)$ for every $F \in \mathbf{M}$. The quotient functor

$$D/R := F \mapsto D(F)/R(F)$$

is then not necessarily k -definably embeddable. In order to remedy this, one simply enlarges our class of functors by *forcing* it to be stable under such quotient operations.

Definition 2.3.3. A functor $\Delta: \mathbf{M} \rightarrow \mathbf{Sets}$ is said to be *k -definable* if it is isomorphic to D/R for some k -definably embeddable functor D and some k -definably embeddable equivalence relation $R \subset D \times D$.

A natural transformation between k -definable functors is called *k -definable* if its graph is itself a k -definable functor.

Remark 2.3.4. Let D be k -definable functor and let $R \subset D \times D$ be a k -definable equivalence relation. It follows from the definition that the quotient functor D/R is k -definable.

(2.3.5) *Comments.* Quantifier elimination ensures that there is no conflict of terminology: if X is a k -scheme of finite type, then a sub-functor D of X is k -definable in the above sense if and only if it is k -definable in the sense of (2.2.3).

The following counterparts of assertions (i) and (ii) of loc. cit. hold (the first one is again a consequence of quantifier elimination; the second one comes from Remark 2.3.4 above).

- *Counterpart of (i).* Let Δ be a k -definable functor. For every $F \in \mathbf{M}$, let us say that a subset of $\Delta(F)$ is k -definable if it can be written $D(F)$ for D a k -definable sub-functor of Δ . The assignment $D \mapsto D(F)$ then induces a one-to-one correspondence between k -definable subsets of $\Delta(F)$ and k -definable sub-functors of Δ . The analogous statement for k -definable transformations holds.
- *Counterpart of (ii).* If $f : D \rightarrow D'$ is a k -definable natural transformation between two k -definable functors, the sub-functor of D' that sends F to the image of $f(F) : D(F) \rightarrow D'(F)$ is k -definable.

Examples 2.3.6. The following functors are k -definable; their descriptions as quotients are left to the readers.

- The functor $\Gamma : F \mapsto |F^\times|$.
- The functor $\Gamma_0 : F \mapsto |F|$.
- The functor $F \mapsto \widetilde{F}$.
- If a and b are elements of $|k|^\mathbb{Q}$ with $a \leq b$, the functor $[a; b]$ that sends F to

$$\{c \in |F|, a \leq c \leq b\}$$

is k -definable (note that $|k|^\mathbb{Q} \subset |F|$ for every $F \in \mathbf{M}$). One defines in an analogous way the functors $[a; b]$, $(a; b)$, $[a; +\infty)$, etc. Such functors are called *k-definable intervals*; a k -definable interval of the form $[a; b]$ will be called a *k-definable segment*.

Remark. One can prove that none of the functors above is k -definably embeddable, except the singleton and the empty set (which can be described as generalized intervals). What it means can be roughly rephrased as follows, say, for Γ_0 (there is an analogous formulation for every of the other functors): one cannot find an algebraic variety X over k and a natural way to embed $|F|$ in $X(F)$ for every $F \in \mathbf{M}$.

(2.3.7) Example of k -definable natural transformations. Let I and J be two k -definable intervals. A *k-monomial map* from I to J is a natural transformation of the form

$$x \mapsto ax^r$$

with $a \in |k|^\mathbb{Q}$ and $r \in \mathbb{Q}_+$ (resp., \mathbb{Q}) if $0 \in I$ (resp., if $0 \notin I$); here we use the convention $0^0 = 1$. Such a natural transformation is k -definable. Moreover, for any k -definable natural transformation $u: I \rightarrow J$, there exists a finite partition $I = \bigsqcup_n I_n$ of I in k -definable intervals such that $u|_{I_n}$ is k -monomial for every n .

Assume now that we are given a k -definable injection $u: I \rightarrow J$, that I contains an interval of the form $]0; a[$ with $a \in |k^\times|^\mathbb{Q}$ and that J is bounded, i.e., contained in $[0; R]$ for some $R \in |k^\times|^\mathbb{Q}$. Then we can always decrease a so that $u|_{]0; a[}$ is k -monomial. Since u is injective, the corresponding exponent will be nonzero; since J is bounded, it will even be positive. Hence the map u induces an increasing k -monomial bijection between $]0; a[$ and a k -definable interval $]0; b[\subset J$.

(2.4) More involved examples of k -definable functors.

(2.4.1) One can *concatenate* a finite sequence of k -definable segments: one makes the quotient of their disjoint union by the identification of successive endpoints and origins; one gets that way k -definable *generalized segments*.

The concatenation of finitely many k -definable segments *with nonzero origin and endpoints* is k -definably isomorphic to a single k -definable segment; the proof consists in writing down an explicit isomorphism and is left to the reader.

But be aware that this does not hold in general without our assumption on the origins and endpoints. For instance, the concatenation I of two copies of $[0; 1]$, where the endpoint 1 of the first one is identified with the origin 0 of the second one, is not k -definably isomorphic to a k -definable segment; indeed, it follows even from (2.3.7) that there is no k -definable injection from I to a bounded k -definable interval.

(2.4.2) A sub-functor of Γ_0^n is k -definable if and only if it can be defined by conjunctions and disjunctions of inequalities of the form

$$ax_1^{e_1} \dots x_n^{e_n} \bowtie bx_1^{f_1} \dots x_n^{f_n}$$

where a and b belong to $|k|$, and where the e_i ’s and the f_i ’s belong to \mathbb{N} , and where $\bowtie \in \{<, \leq, >, \geq\}$.

Any k -definable generalized segment is k -definably isomorphic to some k -definable sub-functor of Γ_0^n (the proof is left to the reader).

(2.4.3) The functor B that sends a field $F \in \mathbf{M}$ to the set of its closed balls is k -definable. Indeed, let us denote by T the k -definable functor

$$F \mapsto \left\{ \begin{pmatrix} a & b \\ 0 & a \end{pmatrix} \right\}_{a \in F^\times, b \in F},$$

and by T' its k -definable sub-functor $F \mapsto T(F) \cap \text{GL}_2(F^\circ)$. The quotient functor T/T' is k -definable, and the reader will check that the map that sends

$$\begin{pmatrix} a & b \\ 0 & a \end{pmatrix}$$

to the closed ball with center b and radius $|a|$ induces a functorial bijection between $T(F)/T'(F)$ and the set of closed balls of F of positive radius. The functor B is thus isomorphic to $F \mapsto (T(F)/T'(F)) \coprod F$, and is therefore k -definable.

(2.5) **About the classification of all k -definable functors.** By definition, the description of a k -definable functor involves a k -definably embeddable functor and a k -definably embeddable equivalence relation on it.

In fact, in the seminal work [14] Haskell et al. prove that the equivalence relation above can always be chosen to be in a particular explicit list, as we now explain.

(2.5.1) Let $n \in \mathbb{N}$. We denote by S_n the functor that sends $F \in \mathbf{M}$ to $\text{GL}_n(F)/\text{GL}_n(F^\circ)$; note that $S_1 \simeq \Gamma$.

We denote by T_n be the functor that sends $F \in \mathbf{M}$ to the quotient of $\text{GL}_n(F)$ by the kernel of $\text{GL}_n(F^\circ) \rightarrow \text{GL}_n(F^\circ/F^{\circ\circ})$.

By construction, S_n and T_n appear as quotients of GL_n by k -definable relations, hence are k -definable.

(2.5.2) Let X be a k -scheme of finite type, and let (n_i) and (m_j) be two finite families of integers. Let D be a k -definable sub-functor of the product $X \times \prod_i \text{GL}_{n_i} \times \prod_j \text{GL}_{m_j}$. The image Δ of D under the map

$$X \times \prod_i \text{GL}_{n_i} \times \prod_j \text{GL}_{m_j} \rightarrow X \times \prod_i S_{n_i} \times \prod_j T_{m_j}$$

is a k -definable sub-functor of $X \times \prod_i S_{n_i} \times \prod_j T_{m_j}$. Haskell et al. have then shown that every k -definable functor is isomorphic to such a Δ .

(2.6) Ground field extension. Let $F \in \mathbf{M}$, let F_0 be a subfield of F containing k , and let $\overline{F_0}$ be the algebraic closure of F_0 inside F . Let \mathbf{N} be the category of all nontrivially valued, algebraically closed fields containing $\overline{F_0}$; note that $F \in \mathbf{N}$, that $\mathbf{N} \subset \mathbf{M}$ (since $\overline{k} \subset \overline{F_0}$), and that $\mathbf{N} = \mathbf{M}$ as soon as F_0 is algebraic over k (since then $\overline{F_0} = \overline{k}$).

The notion of an F_0 -definable functor from \mathbf{N} to **Sets**, and that of an F_0 -definable subset of $\Delta(L)$ for such a functor Δ and for $L \in \mathbf{N}$, makes sense. If D is a k -definable functor from $\mathbf{M} \rightarrow \mathbf{Sets}$, its restriction to \mathbf{N} is F_0 -definable; in particular, the notion of an F_0 -definable subset of $D(L)$ for any $L \in \mathbf{N}$ makes sense.

3 Hrushovski and Loeser's Fundamental Construction

For every $F \in \mathbf{M}$, we denote by \mathbf{M}_F the category of valued extensions of F that belong to \mathbf{M} .

(3.1) The notion of a type (cf. [16, Sects. 2.3–2.9]). Let $F \in \mathbf{M}$ and let D be an F -definable functor.

(3.1.1) Let L and L' be two valued fields belonging to \mathbf{M}_F . Let us say that a point x of $D(L)$ and a point x' of $D(L')$ are *F-equivalent* if for every F -definable subfunctor Δ of D , one has

$$x \in \Delta(L) \iff x' \in \Delta(L').$$

Roughly speaking, x and x' are equivalent if they satisfy the same formulas with parameters in F .

We denote by $S(D)$ the set of couples (L, x) where L belongs to \mathbf{M}_F and where x belongs to $D(L)$, up to F -equivalence; an element of $S(D)$ is called a *type* on D .

(3.1.2) Any $x \in D(F)$ defines a type on D , and one can identify that way $D(F)$ with a subset of $S(D)$, whose elements will be called *simple* types.

(3.1.3) Functoriality. Let f be an F -definable natural transformation from D to another F -definable functor D' and let $x \in S(D)$. If t is a representative of x , then the F -equivalence class of $f(t)$ only depends on x , and not on t ; one thus gets a well-defined type on D' , that will usually be denoted by $f(x)$.

If Δ is an F -definable subfunctor of D , the induced map $S(\Delta) \rightarrow S(D)$ identifies $S(\Delta)$ with the subset of $S(D)$ consisting of types admitting a representative (L, x) with $x \in \Delta(L)$ (it will then be the case for *all* its representatives, by the very definition of F -equivalence); we simply say that such a type *lies on* Δ .

(3.1.4) Let $L \in \mathbf{M}_F$, and let D_L be the restriction of D to \mathbf{M}_L . Let t be a type on D_L . Any representative of t defines a type on D , which only depends on t (indeed, L -equivalence is stronger than F -equivalence); one thus has a natural *restriction* map $S(D_L) \rightarrow S(D)$.

Remark 3.1.5. Let $x \in S(D)$. Let \mathcal{U}_x be the set of subsets of $D(F)$ which are of the kind $\Delta(F)$, where Δ is an F definable sub-functor of D on which x lies. One can rephrase the celebrated compactness theorem of model theory by saying that $x \mapsto \mathcal{U}_x$ establishes a bijection between $S(D)$ and the set of ultra-filters of F -definable subsets of $D(F)$.

(3.1.6) Definable types. Let $t \in S(D)$. By definition, the type t is determined by the data of all F -definable sub-functors of D on which it lies. We will say that t is F -definable if, roughly speaking, the following holds: *for every F -definable family of F -definable sub-functors of D , the set of parameters for which t lies on the corresponding F -definable sub-functor is itself F -definable.*

Let us be more precise. If D' is an F -definable functor, and if Δ is an F -definable sub-functor of $D \times D'$, then for every $L \in \mathbf{M}_F$ and every $x \in D'(L)$, the fiber Δ_x of Δ_L over x is an L -definable sub-functor Δ_x of D_L .

We will say that t is F -definable if for every (D', Δ) as above, the subset of $D'(F)$ consisting of points x such that t lies on Δ_x is F -definable.

(3.1.7) Canonical extension of an F -definable type. If t is an F -definable type on D , then it admits for every $L \in \mathbf{M}_F$ a canonical L -definable pre-image t_L on $S(D_L)$, which is called the *canonical extension* of t .

Roughly speaking, t_L is defined by the same formulas as t . This means the following. Let Σ be an L -definable sub-functor of D_L . Considering the coefficients of the formulas that define Σ as parameters, we see that there exist an F -definable functor D' , an F -definable sub-functor Δ of $D \times D'$, and a point $y \in D'(L)$ such that $\Sigma = \Delta_y$.

Now since t is F -definable, there exists an F -definable sub-functor E of D' such that for every $x \in D'(F)$, the type t lies on Δ_x if and only if $x \in E(F)$. Then the type t_L belongs to $\Sigma = \Delta_y$ if and only if $y \in E(L)$.

(3.1.8) Orthogonality to Γ . Let t be a type on D . We will say that t is *orthogonal to Γ* if it is F -definable and if for every F -definable natural transformation $f : D \rightarrow \Gamma_{0,F}$, the image $f(t)$, which is *a priori* a type on $\Gamma_{0,F}$, is a *simple* type, that is, belongs to $\Gamma_0(F) = |F|$. If this is the case, t_L remains orthogonal to Γ for every $L \in \mathbf{M}_F$.

It follows from the definitions that any simple type on D is orthogonal to Γ .

(3.1.9) Functoriality of those definitions. If $f : D \rightarrow D'$ is an F -definable natural transformation, then for every type $t \in S(D)$ the image $f(t)$ is F -definable (resp., orthogonal to Γ) as soon as t is F -definable (resp., orthogonal to Γ).

(3.2) A very important example: the case of an algebraic variety. Let Y be an algebraic variety over F . The purpose of this paragraph is to get an explicit description of $S(Y)$; as we will see, this descriptions looks like that of the Berkovich analytification of an algebraic variety (1.2.2).

(3.2.1) The general case. Let t be a type on Y . Let (L, x) be a representative of t . The point x of $Y(L)$ induces a *scheme-theoretic* point y of Y and a valuation $|\cdot|_x$ on the residue field $F(y)$, extending that of F ; the data of the point y and of (the equivalence class of) the valuation $|\cdot|_x$ only depend on t , and not on the choice of (L, x) .

Conversely, if y is a schematic point of Y , any valuation on $F(y)$ extending that of F is induced by an isometric embedding $F(y) \hookrightarrow L$ for some $L \in \mathbf{M}_F$, hence arises from some type t on Y . One gets that way a bijection between $S(Y)$ and the *valuative spectrum* of Y , that is, the set of couples $(y, |\cdot|)$ with y a scheme-theoretic point on Y and $|\cdot|$ a valuation on $F(y)$ extending that of F (considered up to equivalence). Note that there is a natural map from $S(Y)$ to Y , sending a couple $(y, |\cdot|)$ to y .

(3.2.2) The affine case. Let us assume now that $Y = \text{Spec } B$, let $y \in Y$ and let $|\cdot|$ be a valuation on $F(y)$ extending that of F . By composition with the evaluation map $f \mapsto f(y)$, we get a valuation $\varphi : B \rightarrow |F(y)|$ (the definition of a valuation on a ring is *mutatis mutandis* the same as on a field—but be aware that it may have a nontrivial kernel); this valuation extends that of F .

Conversely, let φ be a valuation on B extending that of F . Its kernel corresponds to a point y of Y , and φ is the composition of the evaluation map at y and of a valuation $|\cdot|$ on $F(y)$ extending that of F .

These constructions provide a bijection between the valuative spectrum of Y and the set of (equivalence classes of) valuations on B extending that of F ; eventually, we get a bijection between $S(Y)$ and the set of valuations on B extending that of F . Modulo this bijection, the natural map $S(Y) \rightarrow Y$ sends a valuation to its kernel.

(3.2.3) Interpretation of some properties. Let $t \in S(Y)$, let φ be the corresponding valuation on B , and let \mathfrak{p} be the kernel of φ (that is, its image on $Y = \text{Spec } B$). Let \mathfrak{G} be the subset of B^2 consisting of couples (b, b') such that $\varphi(b) \leq \varphi(b')$.

Interpretation of definability. The type t is F -definable if and only if for every finite F -dimensional subspace E of B the following hold:

- the intersection $E \cap \mathfrak{p}$ is an F -definable subset of E ;
- the intersection $(E \times E) \cap \mathfrak{G}$ is an F -definable subset of $E \times E$.

Interpretation of orthogonality to Γ . The type t is orthogonal to Γ if and only if it is F -definable, and if φ takes its values in $|F|$. This is equivalent to require that φ takes its values in $|F|$ and that $\varphi|_E : E \rightarrow |F|$ is F -definable for every finite F -dimensional subspace E of B .

Interpretation of the canonical extension. Assume that the type t is F -definable and let $L \in \mathbf{M}_F$; we want to describe the valuation φ_L on $B \otimes_F L$ that corresponds to the canonical extension t_L (3.1.7). It is equivalent to describe the kernel \mathfrak{p}_L of φ_L and the subset \mathfrak{G}_L of $(B \otimes_F L)^2$ consisting of couples (b, b') such that $\varphi_L(b) \leq \varphi_L(b')$. It is sufficient to describe the intersection \mathfrak{p}_L (resp., \mathfrak{G}_L) with $E \otimes_k L$ (resp., with $(E \otimes_k L)^2$) for any finite dimensional F -vector subspace E of B .

Let us fix such an E . Since t is definable, there exist a unique k -definable sub-functor D of $\underline{E} := \Lambda \mapsto E \otimes_F \Lambda$ such that $E \cap \mathfrak{p} = D(F)$ and a unique k -definable sub-functor D' of \underline{E}^2 such that $(E \times E) \cap \mathfrak{G} = D'(F)$. One then has

$$\mathfrak{p}_L \cap (E \otimes_F L) = D(L) \quad \text{and} \quad \mathfrak{G}_L \cap (E \otimes_F L)^2 = D'(L).$$

Let us assume moreover that t is orthogonal to Γ . In this situation, φ induces a k -definable map $E \rightarrow |F|$. This definable map comes from a unique k -definable transformation $\underline{E} \rightarrow \Gamma_0$. By evaluating it on L , one gets an L -definable map

$$E \otimes_F L \rightarrow |L|,$$

which coincides with the restriction of φ_L .

(3.3) The fundamental definition [16, Sect. 3]. Let V be a k -definable functor. The *stable completion*⁴ of V is the functor

$$\widehat{V} : \mathbf{M} \rightarrow \mathbf{Sets}$$

defined as follows:

- if $F \in \mathbf{M}$, then $\widehat{V}(F)$ is the set of types on V_F that are orthogonal to Γ ;
- if $L \in \mathbf{M}_F$, the arrow $\widehat{V}(F) \rightarrow \widehat{V}(L)$ is the embedding that sends a type t to its canonical extension t_L (3.1.7).

(3.4) Basic properties and first examples.

(3.4.1) The formation of \widehat{V} is functorial in V with respect to k -definable maps (3.1.9).

(3.4.2) Since any simple type (3.1.2) is orthogonal to Γ , one has a natural embedding of functors $V \hookrightarrow \widehat{V}$. We will therefore identify V with a sub-functor of \widehat{V} . For every $F \in \mathbf{M}$, the points of \widehat{V} that belong to $V(F)$ will be called *simple points*.

(3.4.3) The stable completion of a polyhedron. Let $F \in \mathbf{M}$. It follows immediately from the definition that a type on $\Gamma_{0,F}$ is orthogonal to Γ if and only if it is simple. In other words, $\widehat{\Gamma}_0 = \Gamma_0$. This fact extends to k -definable sub-functors of Γ_0^n (2.4.2): if V is such a functor, then the natural embedding $V \hookrightarrow \widehat{V}$ is a bijection.

(3.5) The stable completion of a definable sub-functor of an algebraic variety. Let X be a k -scheme of finite type, and let V be a k -definable sub-functor of X .

(3.5.1) Assume that $V = X = \text{Spec } A$. From (3.2.2) and (3.2.3) we get the following description of \widehat{X} . Let $F \in \mathbf{M}$. The set $\widehat{X}(F)$ is the set of valuations

$$\varphi : A \otimes_k F \rightarrow |F|$$

extending that of F and such that for every finite dimensional F -vector space E of $A \otimes_k F$, the restriction $\varphi|_E : E \rightarrow |F|$ is F -definable. Let $L \in \mathbf{M}_F$ and let $\varphi \in \widehat{X}(F)$. For every finite dimensional F -vector subspace E of $A \otimes_k F$, the F -definable map $\varphi|_E : E \rightarrow |F|$ arises from a unique F -definable natural

⁴This is named after the model-theoretic notion of *stability* which plays a key role in Hrushovski and Loeser’s work (cf. [16, Sects. 2.4, 2.6, and 2.9]), but is far beyond our scope and will thus in some sense remain hidden in this text.

transformation $\Phi_E : \underline{E} \rightarrow \Gamma_0$, which itself gives rise to an L -definable map $\Phi_E(L) : E \otimes_F L \rightarrow |L|$. By gluing the maps $\Phi_E(L)$'s for E going through the set of all finite F -dimensional vector subspaces of A we get a valuation $A \otimes_k L \rightarrow |L|$ extending that of L ; this is precisely the image φ_L of φ under the natural embedding $\widehat{X}(F) \hookrightarrow \widehat{X}(L)$. Roughly speaking, φ_L is “defined by the same formulas as φ .”

We thus see that Hrushovski and Loeser mimic in some sense Berkovich’s construction (1.2.2), but with a model-theoretic and definable flavor.

(3.5.2) Assume that $X = \text{Spec } A$, and that V is defined by a boolean combination of inequalities of the form $|f| \triangleright \lambda |g|$ (with f and g in A and $\lambda \in |k|$). The functor \widehat{V} is then the sub-functor of \widehat{X} consisting of the semi-norms φ satisfying the same combination of inequalities.

(3.5.3) In general, one defines \widehat{V} by performing the above constructions locally and gluing them.

(3.6) A topology on $\widehat{V}(F)$. Let V be a k -definable functor and let $F \in \mathbf{M}$. We are going to define a topology on $\widehat{V}(F)$ in some particular cases; all those constructions will be based upon the *order topology* on $|F|$.

(3.6.1) If V is a k -definable sub-functor of Γ_0^n , then $\widehat{V}(F) = V(F)$ is endowed with the topology induced from the product topology on $\Gamma_0(F) = |F|^n$. In particular if I is a k -definable interval, it inherits a topology.

(3.6.2) If V is a k -definable generalized segment, then $\widehat{V}(F) = V(F)$ inherits a topology using the above construction and the quotient topology.

(3.6.3) If V is a k -scheme of finite type, then $\widehat{V}(F)$ is given the coarsest topology such that:

- for every Zariski-open subset U of V , the subset $\widehat{U}(F)$ of $\widehat{V}(F)$ is open;
- for every affine open subset $U = \text{Spec } A$ of V and every $a \in A \otimes_k F$, the map $\widehat{U}(F) \rightarrow |F|$ obtained by applying valuations to a is continuous.

(3.6.4) If V is a k -definable sub-functor of a k -scheme of finite type X , then $\widehat{V}(F)$ is endowed with the topology induced from that of $\widehat{X}(F)$, defined at (3.6.3) above.

(3.7) Comments.

(3.7.1) Let V be a k -definable functor which is a sub-functor of Γ_0^n , a k -definable generalized segment or a sub-functor of a k -scheme of finite type. For a given $F \in \mathbf{M}$ the set $\widehat{V}(F)$ inherits a topology as described above at (3.6) et sq.

We emphasize that this topology does not depend only on the abstract k -definable functor V , but also on its given presentation. For example, let

$$i: (\text{Spec } k) \coprod \coprod \mathbb{G}_{m,k} \rightarrow \mathbb{A}_k^1$$

be the morphism induced by the closed immersion of the origin and the open immersion of $\mathbb{G}_{m,k}$ into \mathbb{A}_k^1 . The natural transformation induced by i is then a k -definable isomorphism, because for every $F \in \mathbf{M}$ the induced map $\{0\} \coprod F^\times \rightarrow F$ is bijective. But it is not a homeomorphism: $\{0\}$ is open in the left-hand side, and not in the right-hand side.

(3.7.2) Let V be a k -definable sub-functor of some k -scheme of finite type X . The valuation on F defines a topology on $X(F)$; if $V(F)$ is given the induced topology, the inclusion $V(F) \subset \widehat{V}(F)$ is then a *topological embedding with dense image*.

(3.7.3) Let $F \in \mathbf{M}$ and let $L \in \mathbf{M}_F$. Be aware that in general, the topology on $|F|$ induced by the order topology on $|L|$ is *not* the order topology on $|F|$; it is finer. For instance, assume that there exists $\omega \in |L|$ such that $1 < \omega$ and such that there is no element $x \in |F|$ with $1 < x \leq \omega$ (in other words, ω is upper infinitely close to 1 with respect to $|F|$). Then for every $x \in |F^\times|$ the singleton $\{x\}$ is equal to

$$\{y \in |F|, \omega^{-1}x < y < \omega x\}.$$

It is therefore open for the topology induced by the order topology on $|L|$; hence, the latter induces the discrete topology on $|F^\times|$.

Since all the topologies we have considered above ultimately rely on the order topology, this phenomenon also holds for them. That is, let V be as in (3.7.1). The natural embedding $\widehat{V}(F) \hookrightarrow \widehat{V}(L)$ is *not* continuous in general; the topology on $\widehat{V}(F)$ is coarser than the topology induced from that of $\widehat{V}(L)$.

(3.7.4) Let V and W be two k -definable functors, each of which being of one of the form considered in (3.7.1). We will say that a k -definable natural transformation $f : \widehat{V} \rightarrow \widehat{W}$ is *continuous* if $f(F) : \widehat{V}(F) \rightarrow \widehat{W}(F)$ is continuous for every $F \in \mathbf{M}$; we then define in an analogous way the fact for f to be a homeomorphism, or to induce a homeomorphism between \widehat{V} and a sub-functor of \widehat{W} , etc.

(3.8) The affine and projective lines ([16, Example 3.2.1]). Let $F \in \mathbf{M}$.

(3.8.1) For every $a \in F$ and $r \in |F|$, the map $\eta_{a,r,F}$ from $F[T]$ to $|F|$ that sends $\sum a_i(T - a)^i$ to $\max |a_i| \cdot r^i$ is a valuation on $F[T]$ whose associated type belongs to $\widehat{\mathbb{A}}_k^1(F)$. Note that for every $a \in F$, the semi-norm $\eta_{a,0,F}$ is nothing but $P \mapsto |P(a)|$; hence, it corresponds to the simple point $a \in F = \mathbb{A}_k^1(F)$. If $L \in \mathbf{M}_L$, the natural embedding $\widehat{\mathbb{A}}_k^1(F) \hookrightarrow \widehat{\mathbb{A}}_k^1(L)$ is induced by the map that sends the semi-norm $\eta_{a,r,F}$ (for given $a \in F$ and $r \in |F|$) to the semi-norm on $L[T]$ “that is defined by the same formulas,” that is, to $\eta_{a,r,L}$.

Remark 3.8.2. If the ground field F is clear from the context, we will sometimes write simply $\eta_{a,r}$ instead of $\eta_{a,r,F}$.

(3.8.3) For every $(a, b) \in F^2$ and $(r, s) \in |F|^2$, an easy computation shows that the valuations $\eta_{a,r}$ and $\eta_{b,s}$ are equal if and only if $r = s$ and $|a - b| \leq r$, that is, if and only if the closed balls $B(a, r)$ and $B(b, s)$ of F are equal. Moreover, one can prove that every valuation belonging to $\widehat{\mathbb{A}}_k^1(F)$ is of the form $\eta_{a,r}$ for suitable $(a, r) \in F \times |F|$. Therefore we get a functorial *bijection* between $\widehat{\mathbb{A}}_k^1(F)$ and the set of closed balls of F ; it then follows from (2.4.3) that the functor $\widehat{\mathbb{A}}_k^1$ is k -definable.

(3.8.4) The functor $\widehat{\mathbb{P}}_k^1$ is simply obtained by adjoining the simple point ∞ to $\widehat{\mathbb{A}}_k^1$; hence it is k -definable too.

Definitions 3.9. In order to be able to state fundamental results about the stable completions, we need some definitions.

(3.9.1) Let X be a k -scheme of finite type and let V be a k -definable sub-functor of X . Let $F \in \mathbf{M}$. The set of integers n such that there exists an F -definable injection $(F^\circ)^n \hookrightarrow V(F)$ does not depend on F (this follows from quantifier elimination) and is bounded by $\dim X$. Its supremum is called the *dimension of V* ; this is a nonnegative integer $\leq \dim X$ if $V \neq \emptyset$, and $(-\infty)$ otherwise.

Let us mention for further use that there is also a notion of dimension for k -definable sub-functors of Γ_0^n . Let D be such a sub-functor and let $F \in \mathbf{M}$. The set of integers m such that there exists elements $a_1, b_1, \dots, a_m, b_m$ in $|F^\times|$ with $a_i < b_i$ for every i and a k -definable injection

$$[a_1; b_1] \times \dots \times [a_m; b_m] \hookrightarrow D$$

does not depend on F . Its supremum is called the dimension of D . This is a nonnegative integer $\leq n$ if $D \neq \emptyset$, and $(-\infty)$ otherwise.

(3.9.2) A functor from \mathbf{M} to \mathbf{Sets} is said to be *pro- k -definable* if it is isomorphic to a projective limit of k -definable functors. This being said, let us quickly mention two issues we will not discuss in full detail (the interested reader may refer to [16, Sect. 2.2]).

1. If one wants this isomorphism with a projective limit to be canonical, one has to be implicitly given a definition of this functor “in families” [compare to (2.3.1)]; this is the case here: there is a natural notion of a k -definable family of types orthogonal to Γ .
2. The set of indices of the projective system has to be of small enough (infinite) cardinality; in the cases we consider in this paper, it can be taken to be countable.

(3.9.3) Let $V = \text{proj lim } V_i$ and $W = \text{proj lim } W_j$ be two pro- k -definable functors. For every (i, j) , let D_{ij} be the set of all k -definable natural transformations from V_i to W_j . A natural transformation from V to W is said to be *pro- k -definable* if it belongs to

$$\lim_j \left(\lim_i D_{ij} \right).$$

(3.9.4) A functor V from \mathbf{M} to \mathbf{Sets} is said to be *strictly pro- k -definable* if for every pro- k -definable natural transformation f from V to a k -definable functor D , the direct image functor $f(V)$ is a k -definable sub-functor of D (it is *a priori* only a (possibly infinite) intersection of k -definable sub-functors).

(3.9.5) Let V be a pro- k -definable functor from \mathbf{M} to \mathbf{Sets} . A sub-functor D of V is said to be *relatively k -definable* if there exist a k -definable functor Δ , a pro- k -definable natural transformation $f : V \rightarrow \Delta$, and a k -definable sub-functor Δ' of Δ such that $D = f^{-1}(\Delta')$.

Let $F \in \mathbf{M}$. By evaluating those functors at F -points, one gets the notion of a *relatively k -definable* subset of $V(F)$, and more generally of a *relatively F_0 -definable* subset of $V(F)$ for F_0 any valued field lying between k and F [compare to (2.6)].

Replacing F -definable sub-functors by relatively F -definable ones, in the definitions of (3.1) et sq., one gets the notion of a type on V_F , and of an F -definable type on V_F . There is a natural bijection between the set of types on V_F and that of ultra-filters of relatively F -definable subsets of $V(F)$ [compare to Remark 3.1.5].

Theorem 3.10 ([16, Lemma 2.5.1, Theorems 3.1.1, 7.1.1, and Remark 7.1.3]). Let V be a k -definable sub-functor of a k -scheme of finite type. The stable completion \widehat{V} is strictly pro- k -definable, and is k -definable if and only if $\dim V \leq 1$.

(3.10.1) The pro- k -definability of \widehat{V} comes from general arguments of model-theory, which hold in a very general context, that is, not only in the theory of valued fields. Its *strict* pro- k -definability comes from model-theoretic properties of the theory of algebraically closed fields, used at the “residue field” level.

(3.10.2) The definability of \widehat{V} in dimension 1 is far more specific. It ultimately relies on the Riemann–Roch theorem for curves, through the following consequence of the latter: if X is a projective, irreducible, smooth curve of genus g over an algebraically closed field F , the group $F(X)^\times$ is generated by rational functions with at most $g + 1$ poles (counted with multiplicities).

(3.11) Definable compactness. Let V be a k -definable sub-functor of Γ_0^n , a k -definable generalized segment or a k -definable subscheme of a scheme of finite type. The functor \widehat{V} is then pro- k -definable: indeed, in the first two cases it is equal to V , hence even k -definable; in the third case, this is Theorem 3.10 above.

(3.11.1) We have defined at (3.6) et sq. for every $F \in \mathbf{M}$ a topology on $\widehat{V}(F)$. By construction, there exists for every $F \in \mathbf{M}$ a set \mathbb{O}_F of relatively F -definable sub-functors of \widehat{V}_F such that:

- the set $\{D(F)\}_{D \in \mathbb{O}_F}$ is a basis of open subsets of $\widehat{V}(F)$;
- for every $D \in \mathbb{O}_F$ and every $L \in M_F$, the sub-functor D_L of \widehat{V}_L belongs to \mathbb{O}_L (in particular, $D(L)$ is an open subset of $\widehat{V}(L)$).

Let $F \in \mathbf{M}$ and let t be a type on \widehat{V}_F (this makes sense in view of (3.9.5) since \widehat{V} is pro- k -definable by Theorem 3.10). We say that a point $x \in \widehat{V}(F)$ is a *limit* of t if the following holds: *for every $D \in \mathbb{O}_F$ such that $x \in D(F)$, the type t lies on D* . If X is separated, $\widehat{V}(F)$ is Hausdorff, and the limit of a type is then unique provided it exists.

(3.11.2) We will say that \widehat{V} is *definably compact* if for every $F \in \mathbf{M}$, every F -definable type on \widehat{V}_F has a unique limit in $\widehat{V}(F)$.

Remarks 3.11.3. Viewing types as ultra-filters, we may rephrase the above definition by saying that \widehat{V} is definably compact if for every $F \in \mathbf{M}$, every F -definable ultra-filter of relatively F -definable subsets of $\widehat{V}(F)$ has a unique limit on $\widehat{V}(F)$.

(3.12) The following propositions provide evidence for definable compactness being the right model-theoretic analogue of usual compactness; they are particular cases of some general results proved (or at least stated with references) in [16, Sects. 4.1 and 4.2].

Proposition 3.12.1. *Let V be a k -definable sub-functor of Γ_0^n . If V is contained in $[0; R]^n$ for some $R \in |k|^\mathbb{Q}$ and can be defined by conjunction and disjunction of non-strict monomial inequalities, then V is definably compact.*

Remark 3.12.2. Any k -definable generalized segment is definably compact: this can be proved either directly, or by writing down an explicit k -definable homeomorphism between such a segment and a functor V as in Proposition 3.12.1 above.

Proposition 3.12.3. *Let X be a k -scheme of finite type.*

- (1) *The stable completion \widehat{X} is definably compact if and only if X is proper.*
- (2) *Assume that X is proper; let (X_i) be an affine covering of X and let V be a k -definable sub-functor of X such that $V \cap X_i$ can be defined for every i by conjunction and disjunction of non-strict inequalities. The stable completion \widehat{V} is then definably compact.*

4 Homotopy Type of \widehat{V} and Links with Berkovich Spaces

(4.1) The link between stable completions and Berkovich spaces. For this paragraph, we assume that the valuation $|\cdot|$ of k takes real values, and that k is complete.

(4.1.1) Let X be an algebraic variety over k , and let V be a k -definable sub-functor of X . The inequalities that define V also define a strict semi-algebraic subset V^{an} of X^{an} [for the definition of such a subset see (1.3.3)]; as the notation suggests, V^{an} only depends on the sub-functor V of X , and not on its chosen description (this follows from quantifier elimination). Any strict semi-algebraic subset of X^{an} is of that kind.

(4.1.2) Let $F \in \mathbf{M}$ be such that $|F| \subset \mathbb{R}_+$. Any point of $\widehat{V}(F)$ can be interpreted in a suitable affine chart as a valuation with values in $|F| \subset \mathbb{R}_+$; it thus induces a point of V^{an} . We get that way a map $\pi : \widehat{V}(F) \rightarrow V^{\text{an}}$ which is continuous by the very definitions of the topologies involved. One can say a bit more about π in some particular cases; the following facts are proved in [16, Sect. 14].

(4.1.3) *We assume that k is algebraically closed, nontrivially valued and that $F = k$. The map π then induces a homeomorphism onto its image.*

Let us describe this image for $V = X = \mathbb{A}_k^1$. By the explicit description of $\widehat{\mathbb{A}}_k^1$, it consists precisely of the set of points $\eta_{a,r}$ with $a \in k$ and $r \in |k|$, that is, of the set of points of type 1 and 2 according to Berkovich's classification (see [2, Chaps. 2 and 4]).

This generalizes as follows: if X is any curve, then $\pi(\widehat{V}(k))$ is precisely the set of points of V^{an} that are of type 1 or 2.

Comments. The fact that Hrushovski and Loeser’s theory, which focuses on definability, only sees points of type 1 and type 2, encodes the following (quite vague) phenomenon: *when one deals only with algebraic curves and scalars belonging to the value groups of the ground field, points of type 1 and type 2 are the only ones at which something interesting may happen.*

Of course, considering all Berkovich points is useful because it ensures good topological properties (like pathwise connectedness and compactness if one starts from a projective, connected variety). Those properties are lost when one only considers points of type 1 and type 2; Hrushovski and Loeser remedy it by introducing the corresponding model-theoretic properties, like definable compactness, or definable arcwise connectedness (see below for the latter).

(4.1.4) *We assume that $|F| = \mathbb{R}_+$ and that F is maximally complete (i.e., it does not admit any proper valued extension with the same value group and the same residue field).* In this case (without any particular assumption on k) the continuous map π is a proper surjection.

(4.1.5) *We assume that $|k| = \mathbb{R}_+$, that k is algebraically closed and maximally complete, and that $F = k$.* Fitting together (4.1.3) and (4.1.4) we see that π establishes then a homeomorphism $\widehat{V}(k) \simeq V^{\text{an}}$.

Suppose for the sake of simplicity that $V = X = \text{Spec } A$. By definition, a point of X^{an} is a valuation $\varphi : A \rightarrow \mathbb{R}_+ = |k|$ extending that of k . Therefore the latter assertion says that for any finite dimensional k -vector space E of A the restriction $\varphi|_E$ is *automatically* k -definable, because of the maximal completeness of k . This “automatic definability” result comes from the previous work [15] by Haskell, Hrushovski, and Macpherson about the model theory of valued fields (which is intensively used by Hrushovski and Loeser throughout their paper); the reader who would like to have a more effective version of it with a somehow explicit description of the formulas describing $\varphi|_E$ may refer to the recent work [18] by Poineau.

(4.2) Hrushovski and Loeser therefore work *purely inside the “hat” world*; they transfer thereafter their results to the Berkovich setting (in [16, Sect. 14]) using the aforementioned map π and its good properties in the maximally complete case (4.1.4) and (4.1.5).

For that reason, *we will not deal with Berkovich spaces anymore*. We are now going to state the “hat-world” avatar of (1.4.1) (this is Theorem 4.3 below) and sketch its proof very roughly.

In fact, Theorem 4.3 is more precisely the “hat world” avatar of a *weakened* version of (1.4.1), consisting of the same statement for *strict* (instead of general) semi-algebraic subsets of Berkovich spaces. But this is not too serious a restriction: indeed, Hrushovski and Loeser prove their avatar of (1.4.1) by first reducing it to the strict case through a nice trick, consisting more or less in “seeing bad scalars as extra parameters” ([16], beginning of Sect. 11.2; the reader may also refer to [12, Sect. 4.3] for detailed explanations).

Theorem 4.3 (Simplified Version of Theorem 11.1.1 of Hrushovski and Loeser [16]). *Let X be a quasi-projective k -variety, and let V be a k -definable sub-functor of X . Let G be a finite group acting on X and stabilizing V , and let E be a finite set of k -definable transformations from V to Γ_0 (if $f \in E$, we still denote by f the induced natural transformation $\widehat{V} \rightarrow \widehat{\Gamma}_0 = \Gamma_0$; if $g \in G$, we still denote by g the induced automorphism of \widehat{V}). There exist:*

- a k -definable generalized segment I , with endpoints o and e ;
- a k -definable sub-functor S of \widehat{V} (where S stands for skeleton);
- a k -definable sub-functor P of Γ_0^m (for some m) of dimension $\leq \dim X$;
- a k' -definable homeomorphism $S \simeq P$, for a suitable finite extension k' of k inside \bar{k} ;
- a continuous k -definable map $h : I \times \widehat{V} \rightarrow S$ satisfying the following properties for every $F \in \mathbf{M}$, every $x \in \widehat{V}(F)$, every $t \in I(F)$, every $g \in G$, and every $f \in E$:
 - $h(o, x) = x$, and $h(e, x) \in S(F)$;
 - $h(t, x) = x$ if $x \in S(F)$;
 - $h(e, h(t, x)) = h(e, x)$;
 - $f(h(t, x)) = f(x)$;
 - $g(h(t, x)) = h(t, g(x))$.

(4.4) Comments.

(4.4.1) We will refer to the existence of k' , of P , and of the k' -definable homeomorphism $S \simeq P$ by saying that S is a *tropical subfunctor* of \widehat{V} . The finite extension k' of k cannot be avoided; indeed, it reflects the fact that the Galois action on the homotopy type of \widehat{V} is not necessarily trivial. In the Berkovich language, think of the \mathbb{Q}_3 -elliptic curve $E : y^2 = x(x-1)(x-3)$. The analytic curve $E_{\mathbb{Q}_3(i)}^{\text{an}}$ admits a Galois-equivariant deformation retraction to a circle, on which the conjugation exchanges two half-circles; it descends to a deformation retraction of E^{an} to a compact interval.

(4.4.2) We will sum up the three first properties satisfied by h by simply calling h a *homotopy with image S* , and the two last ones by saying that h *preserves the functions belonging to E and commutes with the elements of G* .

(4.4.3) The quasi-projectivity assumption can likely be removed, but it is currently needed in the proof for technical reasons.

(4.4.4) The theorem does not simply assert the existence of a (model-theoretic) deformation retraction of \widehat{V} onto a tropical sub-functor; it also ensures that this deformation retraction can be required to preserve finitely many arbitrary natural transformations from V to Γ_0 , and to commute with an arbitrary algebraic action of a finite group. This strengthening of the expected statement is of course intrinsically interesting, but this is not the only reason why Hrushovski and Loeser have decided to prove it. Indeed, even if one only wants to show the existence of a deformation retraction onto a tropical functor with no extra requirements, there is a crucial step in the proof (by induction) at which one needs to exhibit a deformation retraction of a lower dimensional space to a tropical functor *which preserves some natural transformations to Γ_0 and the action of a suitable finite group*.

(4.4.5) The remaining part of this section will now be devoted to (a sketch of) the proof of Theorem 4.3. We will first explain (4.5)–(4.7) what happens for curves; details can be found in [16, Sect. 7].

(4.5) The stable completion $\widehat{\mathbb{P}}_k^1$ is uniquely path-connected.

(4.5.1) Let $F \in \mathbf{M}$ and let x and y be two points of $\widehat{\mathbb{P}}_k^1(F)$. One proves the following statement, which one can roughly rephrase by saying that $\widehat{\mathbb{P}}_k^1$ is uniquely path-connected, or is a tree: *there exists a unique F -definable sub-functor $[x; y]$ of $\widehat{\mathbb{P}}_F^1$ homeomorphic to a k -definable generalized segment with endpoints x and y .*

(4.5.2) Let us describe $[x; y]$ explicitly. If $x = y$, there is nothing to do; if not, let us distinguish two cases. In each of them, we will exhibit an F -definable generalized interval I and an F -definable natural transformation $\varphi : I \rightarrow \widehat{\mathbb{P}}_F^1$ inducing a homeomorphism between I and a sub-functor of $\widehat{\mathbb{P}}_F^1$, and sending one of the endpoints of I to x , and the other one to y .

(4.5.3) *The case where, say, $y = \infty$.* One then has $x = \eta_{a,r}$ for some $a \in F$ and some $r \in |F|$. We take for I the generalized interval $[r; +\infty]$ defined in a natural way: if $r > 0$, it is homeomorphic to the interval $[0; 1/r]$; if $r = 0$, one concatenates $[0; 1]$ and $[1; +\infty]$. We take for φ the natural transformation given by the formula $s \mapsto \eta_{a,s}$, with the convention that $\eta_{a,+\infty} = \infty$.

(4.5.4) *The case where $x = \eta_{a,r}$ and $y = \eta_{b,s}$ for $a, b \in F$ and $r, s \in |F|$.* We may assume that $r \leq s$.

- If $|a - b| \leq s$, then $y = \eta_{a,s}$. We take for I the F -definable interval $[r; s]$ and for φ the natural transformation given by the formula $t \mapsto \eta_{a,t}$.
- If $|a - b| > s$, we take for I the concatenation of $I_1 := [r; |a - b|]$ and of $I_2 := [|a - b|; s]$, and for φ the natural transformation given by the formulas $t \mapsto \eta_{a,t}$ for $t \in I_1$ and $t \mapsto \eta_{b,t}$ for $t \in I_2$.

(4.5.5) If Δ is a finite, k -definable (that is, Galois-invariant) subset of $\widehat{\mathbb{P}}_k^1$, the sub-functor $\bigcup_{x,y \in \Delta} [x; y]$ of $\widehat{\mathbb{P}}_k^1$ is called the *convex hull* of Δ . It is a k -definable “twisted finite sub-tree” of $\widehat{\mathbb{P}}_k^1$, where “twisted” refers to the fact that the Galois action on this sub-tree is possibly nontrivial.

(4.6) Retractions from $\widehat{\mathbb{P}}_k^1$ to twisted finite sub-trees [16, Sect. 7.5].

(4.6.1) Let U and V be the k -definable sub-functors of $\widehat{\mathbb{P}}_k^1$, respectively described by the conditions $|T| \leq 1$ and $|T| \geq 1$. Note that V is k -definably isomorphic to U through $\psi : T \mapsto 1/T$; we still denote by ψ the induced isomorphism $\widehat{V} \simeq \widehat{U}$.

Let $F \in \mathbf{M}$ and let $x \in \widehat{U}(F)$. One has $x = \eta_{a,r}$ for some $a \in F^\circ$ and some $r \in |F^\circ|$. For every $t \in |F^\circ|$, set

$$h(t, x) = \eta_{a, \max(t, r)},$$

and for every $y \in \widehat{V}(F)$, set

$$h(t, y) = \psi^{-1}(h(t, \psi(y))).$$

One immediately checks that both definitions of h agree on $\widehat{U} \cap \widehat{V}$, and we get by gluing a homotopy $h : [0; 1] \times \widehat{\mathbb{P}}_k^1 \rightarrow \widehat{\mathbb{P}}_k^1$ with image $\{\eta_{0,1}\}$. Hence $\widehat{\mathbb{P}}_k^1$ is “ k -definably contractible.”

(4.6.2) Let Δ be a finite, k -definable subset of $\widehat{\mathbb{P}}_k^1$, and let D be the convex hull of $\Delta \cup \{\eta_{0,1}\}$ (4.5.5). Define $h_\Delta : [0; 1] \times \widehat{\mathbb{P}}_k^1 \rightarrow \widehat{\mathbb{P}}_k^1$ as follows: for every x , we denote by τ_x the smallest time t such that $h(t, x) \in D$, and we set $h_\Delta(t, x) = h(t, x)$ if $t \leq \tau_x$, and $h_\Delta(t, x) = h(\tau_x, x)$ otherwise. The natural transformation h_Δ is then a homotopy, whose image is the twisted polyhedron D .

(4.7) Retractions from a curve to a twisted finite graph. In this paragraph we will sketch the proof of the theorem when $X = V$ is a projective algebraic curve; the reader will find more details in the survey [12, Sect. 4.2]. One first chooses a finite G -equivariant map $f : X \rightarrow \mathbb{P}_k^1$, inducing a natural transformation $\widehat{f} : \widehat{X} \rightarrow \widehat{\mathbb{P}}_k^1$.

(4.7.1) The key point is the following [16, Theorem 7.5.1]: *there exists a finite k -definable subset Δ_0 of $\mathbb{P}^1(\bar{k})$ (or, in other words, a divisor on \mathbb{P}_k^1) such that for every finite k -definable subset Δ of $\mathbb{P}^1(\bar{k})$ containing Δ_0 , the homotopy h_Δ lifts uniquely to a homotopy $h_\Delta^X : [0; 1] \times \widehat{X} \rightarrow \widehat{X}$.* Hrushovski and Loeser prove it by carefully analyzing the behavior of the cardinality of the fibers \widehat{f} , as a function from $\widehat{\mathbb{P}}_k^1$ to \mathbb{N} [16, Proposition 7.4.5]. The definability of \widehat{X} and $\widehat{\mathbb{P}}_k^1$ plays a crucial role for that purpose.

(4.7.2) Let us now explain why (4.7.1) allows to conclude. Let Δ be as above, let D be the convex hull of $\Delta \cup \{\eta_{0,1}\}$, and set $D' = \widehat{f}^{-1}(D)$.

- By a definability argument, D' is a “twisted finite sub-graph” of \widehat{X} .
- By choosing Δ sufficiently big, one may ensure that every function belonging to E is locally constant outside D' . This comes from the fact that every k -definable natural transformation from \widehat{X} to Γ_0 is locally constant outside a twisted finite subgraph of \widehat{X} : indeed, any k -definable function can be described by using only norms of regular functions; and the result for such a norm is deduced straightforwardly from the behavior of $|T|$ on $\widehat{\mathbb{P}}_k^1$, which is locally constant outside $[0; \infty]$.

This implies that h_Δ^X preserves the functions belonging to E . Moreover the uniqueness of h_Δ^X ensures that it commutes with the elements of G , and we are done.

(4.7.3) A consequence: path-connectedness of stable completions. Let W be a k -definable sub-functor of a k -scheme of finite type, a k -definable sub-functor of Γ_0^n or a k -generalized interval. We say that \widehat{W} is *definably path-connected*

if for every $F \in \mathbb{M}$ and every $(x, y) \in \widehat{W}(F)^2$ there exists an F -definable generalized interval I with endpoints o and e , and an F -definable continuous natural transformation $h : I \rightarrow \widehat{W}_F$ such that $h(o) = x$ and $h(e) = y$.

It follows easily from the above that if W is a projective curve, then \widehat{W} has finitely many path-connected components, each of which is \bar{k} -definable (indeed, this holds for any twisted finite subgraph of \widehat{W}). Starting from this result, Hrushovski and Loeser prove in fact that if Y is a k -algebraic variety, \widehat{Y} is definably path-connected as soon as Y is geometrically connected [16, Theorem 10.4.2]; in general, there is a Galois-equivariant bijection between the set of geometrical connected components of Y and that of path-connected components of \widehat{Y} .

(4.8) The general case of Theorem 4.3: preliminaries.

(4.8.1) First reductions. By elementary geometrical arguments ([16], beginning of Sect. 11.2) one proves that there exists a G -equivariant embedding of X as a subscheme of an equidimensional projective k -variety X^\sharp on which G acts. Hence by replacing X with X^\sharp we may assume that X is projective and of pure dimension n for some n ; by extending the scalars to the perfect closure of k and by replacing X with its underlying reduced subscheme, one can also assume that its smooth locus is dense. *One proceeds then by induction on n .* The case $n = 0$ is obvious; now one assumes that $n > 0$ and that the theorem holds for smaller integers.

We will explain how to build a homotopy from the *whole space* \widehat{X} to a tropical sub-functor S of \widehat{X} , which commutes with the elements of G and preserves the characteristic function of V and the functions belonging to E ; this homotopy will stabilize \widehat{V} and the image of \widehat{V} will be a k -definable sub-functor of S (since V is *strictly pro-definable*), hence will be tropical as well.

By adding the characteristic function of V to E we thus reduce to the case where $V = X$.

(4.8.2) Use of blowing-up to get a curve fibration. One can blow up X along a *finite set of closed points* so that the resulting variety X' admits a morphism $X' \rightarrow \mathbb{P}_k^{n-1}$ whose generic fiber is a curve; *this is the point where we need X to be projective.* By proceeding carefully (and paying a special attention to the case where k is finite) one can even ensure the following [16, Sect. 11.2]:

- The action of G on X extends to X' , and $X' \rightarrow \mathbb{P}_k^{n-1}$ is G -equivariant.
- There exists a G -equivariant divisor D_0 on X' , finite over \mathbb{P}_k^{n-1} and containing the exceptional divisor D of $X' \rightarrow X$, and a G -equivariant étale map from $X' \setminus D_0$ to \mathbb{A}_k^n (in particular, D_0 contains the singular locus of X').
- There exists a nonempty open subset U of \mathbb{P}_k^{n-1} , whose pre-image on X' is the complement of a divisor D_1 , and a factorization of $X' \setminus D_1 \rightarrow U$ through a finite, G -equivariant map $f : X' \setminus D_1 \rightarrow \mathbb{P}_k^1 \times_k U$.

If one builds a homotopy from \widehat{X}' to a twisted polyhedron which commutes with the action of G , preserves the functions belonging to the (pullback of) E , and the characteristic function of the exceptional divisor D , it will descend to a homotopy

on \widehat{X} satisfying the required properties, because every connected component of \widehat{D} collapses to a point. Hence we reduce to the case where $X' = X$.

(4.9) The concatenation of a homotopy on the base and a fiberwise homotopy. This step is the core of the proof; we sum up here what is done in [16, Sects. 11.3 and 11.4].

(4.9.1) One first applies the *relative* version of the general construction we have described in (4.7) to the finite map f . It provides a divisor Δ on $\mathbb{P}_k^1 \times_k U$ finite over U and a “fiberwise” homotopy h_Δ on $\widehat{\mathbb{P}_k^1 \times_k U}$ which lifts uniquely to a homotopy h_Δ^X on $X \setminus \widehat{D}_1$ whose image is a *relative (at most) one-dimensional tropical functor* over \widehat{U} . By taking Δ big enough, one ensures that h_Δ^X preserves the functions belonging to E , and commutes with the elements of G . Moreover, we can also assume that the pre-image of Δ on $X \setminus D_1$ contains $D_0 \setminus D_1$. Under this last assumption, $\widehat{D_0 \setminus D_1}$ is pointwise fixed under h_Δ^X at every time; therefore, h_Δ^X extends to a homotopy (which is still denoted by h_Δ^X) on $X \setminus \widehat{D}_1 \cup \widehat{D}_0$ which fixes \widehat{D}_0 pointwise at every time. This homotopy preserves the functions belonging to E and commutes with the action of G , and its image Υ is a relative (at most) one-dimensional tropical functor over $\widehat{\mathbb{P}_k^{n-1}}$. Indeed, this is true by construction over \widehat{U} ; and over the complement of \widehat{U} , the fibers of Υ coincide with those of \widehat{D}_0 , which are finite.

(4.9.2) Use of the induction hypothesis. Our purpose is now to exhibit a homotopy h^Υ on Υ , preserving the functions belonging to E and commuting with the action of G , whose image is tropical of dimension at most n . The key point allowing such a construction is a theorem by Hrushovski and Loeser [16, Theorem 6.4.4] which ensures that this problem is under algebraic control, in the following sense: there exists a finite quasi-Galois cover $Z \rightarrow \mathbb{P}_k^{n-1}$, with Galois group H , and a finite family F of k -definable natural transformation from Z to Γ_0 , such that for every homotopy λ on $\widehat{\mathbb{P}_k^{n-1}}$ the following are equivalent:

- (i) λ lifts to a homotopy on \widehat{Z} preserving the functions belonging to F ;
- (ii) λ lifts to a homotopy on Υ preserving the functions belonging to E and commuting with the action of G .

Now *by induction hypothesis* there exists a homotopy λ' on \widehat{Z} , preserving the functions belonging to F and commuting with the action of H , and whose image is tropical of dimension $\leq n - 1$. Since λ' commutes with the action of H , it descends to a homotopy λ on $\widehat{\mathbb{P}_k^{n-1}}$. By its very definition, λ satisfies condition (i) above, hence also condition (ii). Let h^Υ be a homotopy on Υ lifting λ , preserving the functions belonging to E , and commuting with the action of G . Since the image of λ' is tropical of dimension $\leq n - 1$, so is the image of λ (because it is the quotient of that of λ' by the action of the finite group H). As a consequence, the image Σ of h^Υ is a relative (at most) one-dimensional tropical functor over a tropical functor of dimension $\leq n - 1$; therefore, Σ is itself tropical of dimension $\leq n$.

By making h^Υ follow h_Δ^X , we get a homotopy h^0 on $\widehat{X \setminus D_1} \cup \widehat{D_0}$, preserving the functions belonging to E and commuting with the action of G , whose image is the tropical functor Σ .

(4.10) Fleeing away from $\widehat{D_1}$.

(4.10.1) The inflation homotopy: quick description. Hrushovski and Loeser define [16, Lemma 10.3.2] a “inflation” homotopy $h^{\text{inf}} : [0; 1] \times \widehat{X} \rightarrow \widehat{X}$ which fixes pointwise $\widehat{D_0}$ at every time, preserves the functions belonging to E and commutes with the action of G , and which is such that $h^{\text{inf}}(t, x) \in \widehat{X \setminus D_1}$ for every $x \notin \widehat{D_0}$ and every $t > 0$.

(4.10.2) The inflation homotopy: construction. One first defines a homotopy α on $\widehat{\mathbb{A}_k^n}$ by “making the radii of balls increase” (or, in other words, by generalizing to the higher dimension case the formulas we have given for the affine line). It has the following genericity property: if $x \in \widehat{\mathbb{A}_k^n}(F)$ for some $F \in \mathbb{M}$, if Y is a $(n - 1)$ -dimensional Zariski-closed subset of \mathbb{A}_F^n , and if t is a nonzero element of $|F^\circ|$, then $\alpha(t, x) \notin \widehat{Y}(F)$. The homotopy α is lifted to a homotopy μ on $\widehat{X \setminus D_0}$ thanks to the étale G -equivariant map $X \setminus D_0 \rightarrow \mathbb{A}_k^n$, and the genericity property of α is transferred to μ . The homotopy h^{inf} is then defined using a suitable “stopping time function” $x \mapsto \tau_x$ from $\widehat{X \setminus D_0}$ to Γ by the formulas:

- $h^{\text{inf}}(t, x) = \mu(t, x)$ if $x \notin \widehat{D_0}$ and if $t \leq \tau_x$;
- $h^{\text{inf}}(t, x) = \mu(\tau_x, x)$ if $x \notin \widehat{D_0}$ and if $t \geq \tau_x$;
- $h^{\text{inf}}(t, x) = x$ if $x \in \widehat{D_0}$.

(4.10.3) Let us quickly explain why h^{inf} satisfies the required properties.

- Its continuity comes from the choice of the stopping time τ_x : the closer x is to $\widehat{D_0}$, the smaller is τ_x .
- The G -equivariance of h^{inf} comes from its construction, and from the G -equivariance of $X \setminus D_0 \rightarrow \mathbb{A}_k^n$.
- The fact that $h^{\text{inf}}(t, x) \in \widehat{X \setminus D_1}$ for every $x \notin \widehat{D_0}$ and every $t > 0$ is a particular case of the aforementioned genericity property of μ .
- ◊ The fact that h^{inf} preserves the functions belonging to E goes as follows. Those functions are defined using norms of regular functions; now if a regular function is invertible at a point x of $\widehat{X \setminus D_0}$, then its norm is constant in a neighborhood of x , hence will be preserved by $h^{\text{inf}}(\cdot, x)$ if τ_x is small enough; this is also obviously true, if the function vanishes in the neighborhood of x .

Of course, it may happen that the zero locus of a regular function involved in the description of E has some $(n - 1)$ -dimensional irreducible component Y , and if $x \in \widehat{Y \setminus D_0}$, none of the above arguments will apply around it. One overcomes this forthcoming issue by an additional work at the very beginning of the proof: one simply includes such bad components in the divisor D_0 .

(4.10.4) By first applying h^{inf} , and then h^0 , one gets a continuous natural transformation $h^1 : J \times \widehat{X} \rightarrow \widehat{X}$ for some k -definable generalized interval J . Its image is a k -definable sub-functor Σ' of Σ , hence is in particular tropical of dimension $\leq n$.

(4.11) The tropical homotopy and the end of the proof.

(4.11.1) The homotopy h^1 of (4.10.4) now enjoys all required properties, except (possibly) one of them: there is no reason why Σ' should be pointwise fixed at every time, because h^{inf} could disturb it.

(4.11.2) Hrushovski and Loeser remedy it as follows. They proceed to the above construction in such a way that there exists a k -definable G -equivariant sub-functor Σ_0 of Σ satisfying the following conditions.

- (1) Σ_0 is pointwise fixed at every time under h^1 . This is achieved by ensuring that Σ_0 is purely n -dimensional and that there exists a finite extension k' of k and a k' -definable continuous injection $\Sigma_0 \hookrightarrow \Gamma_0^m$ that is preserved by h^1 ; this implies the required assertion by Hrushovski and Loeser [16, Proposition 8.3.1].
- (2) There exists a homotopy h^{trop} on Σ , preserving the functions belonging to E and commuting with the action of G , and whose image is precisely Σ_0 . The construction is purely tropical and rather technical, and we will not give any detail here (see [16, Sect. 11.5]).

(4.11.3) The expected homotopy h on \widehat{X} is now defined by applying first h^1 , and then h^{trop} . Its image S is equal to $h^{\text{trop}}(\Sigma')$, hence is a k -definable sub-functor of Σ_0 . By condition 1), it is pointwise fixed at every time under h^1 , hence under h .

5 An Application of the Definability of \widehat{C} for C a Curve

(5.1) **Presentation of the main result.** We fix a field k which is complete with respect to a non-Archimedean absolute value.

(5.1.1) Let us introduce some vocabulary (essentially coming from [3, Sect. 1.3]). If Z is a topological space and if Z' is a subset of Z , a family (Z_i) of subsets of Z' is said to be a G -covering of Z' (where G stands for ‘‘Grothendieck’’) if every point z of Z' has a neighborhood in Z' of the form $\bigcup_{i \in J} Z_i$ where J is a finite set of indices such that $z \in Z_i$ for every $i \in J$.

Any k -analytic space X is equipped with a class of distinguished compact subsets, the *affinoid domains*; the typical example to have in mind is the subset of $\mathbb{A}_k^{n, \text{an}}$ defined by the inequalities

$$|T_1| \leq r_1 \text{ and } \dots \text{ and } |T_n| \leq r_n$$

where the r_i 's are positive real numbers (this is the ‘‘Berkovich compact polydisc of polyradius (r_1, \dots, r_n) ’’).

The subsets of X that are G -covered by the affinoid domains they contained are called *analytic domains*, and inherit a canonical k -analytic structure. The open subsets of X are analytic domains; so are the compact subsets of X which are unions of finitely many affinoid domains.

(5.1.2) About piecewise-linear geometry (after Berkovich, [6, Sect. 1]). We will follow Berkovich’s convention in piecewise-linear geometry, which is based upon the *multiplicative notation*, in order to avoid using many somehow irrelevant “log” symbols. We will thus say that a subset of $(\mathbb{R}_+^\times)^n$ is a polyhedron if its image under \log is a polyhedron of \mathbb{R}^n (Definition 1.2.5), and that a map between two polyhedra is piecewise-linear if it is piecewise \mathbb{Q} -linear in the usual sense modulo \log isomorphisms.

A *polyhedral structure* on a compact topological space P is a map $P \rightarrow (\mathbb{R}_+^\times)^n$ (for some n) inducing a homeomorphism between P and a polyhedron of $(\mathbb{R}_+^\times)^n$. Two polyhedral structures $i: P \rightarrow (\mathbb{R}_+^\times)^n$ and $j: P \rightarrow (\mathbb{R}_+^\times)^m$ are said to be *equivalent* if there exists a PL homeomorphism $u: i(P) \simeq j(P)$ such that $j = u \circ i$.

A *polyhedral chart* on a topological space Z is a compact subset P of Z (the *support* of the chart) equipped with a class of polyhedral structures. Two polyhedral charts $i: P \simeq P_0$ and $j: Q \simeq Q_0$ are *compatible* if both $i(P \cap Q)$ and $j(P \cap Q)$ are polyhedra (this depends only of the classes of i and j).

A *polyhedral atlas* on a topological space S is a family of polyhedral charts which are pairwise compatible and whose support G -cover S . Two polyhedral atlases on S are said to be *equivalent* if their union is a polyhedral atlas, that is, if any chart of the first one is compatible with any chart of the second one. A *piecewise-linear space*, or PL space for short, is a topological space equipped with an equivalence class of polyhedral atlases, or with a maximal polyhedral atlas—this amounts to the same. When we will speak about a polyhedral atlas on a given PL space, we will implicitly assume that it belongs to the equivalence class defining the PL structure involved.

If S is a PL space, a *PL subspace* of S is a subset of S which is G -covered by (support of) charts of the maximal polyhedral atlas of S ; such a subspace Σ inherits a canonical PL structure. Open subsets of S and compact subsets which are the union of finitely many (support of) charts of the maximal polyhedral atlas of S are examples of PL subspaces of S .

A map $f: S \rightarrow \Sigma$ between PL spaces is said to be PL if there is a polyhedral atlas \mathfrak{A} on S such that for every $P \in \mathfrak{A}$ the image $f(P)$ belongs to the maximal polyhedral atlas of Σ , and $P \rightarrow f(P)$ is PL modulo the identifications of its source and its target with polyhedra. If moreover \mathfrak{A} can be chosen so that $P \rightarrow f(P)$ is injective for every $P \in \mathfrak{A}$, we say that f is a *piecewise immersion*.

(5.1.3) Let $n \in \mathbb{N}$. Once equipped with the family of all its polyhedra, $(\mathbb{R}_+^\times)^n$ becomes a PL space.

Let us denote by S_n the “skeleton” of $\mathbb{G}_m^{n,\text{an}}$. This is the set of semi-norms of the form $\eta_{r_1, \dots, r_n} := \sum a_j T^j \mapsto \max |a_j| r^j$. The very definition of S_n provides it with a homeomorphism onto $(\mathbb{R}_+^\times)^n$, through which it inherits a piecewise-linear structure.

(5.1.4) Let X be a k -analytic space of dimension n , and let $f : X \rightarrow \mathbb{G}_m^{n,\text{an}}$ be a morphism. In [10] (see also the *erratum* [11]) the author has proven that $f^{-1}(S_n)$ inherits a canonical PL structure, with respect to which $f^{-1}(S_n) \rightarrow S_n$ is a piecewise immersion. This generalizes his preceding work [8], in which some additional assumptions on X were needed, and in which the canonicity of the PL structure (answering a question by Temkin) had not been addressed.

Moreover, the proof given in [8] used de Jong’s alterations, while that given in [10] replaces it by the k -definability of \widehat{C} for C an algebraic curve (Theorem 3.10) which simply comes from Riemann–Roch theorem for curves— together with some general results about model theory of valued fields that are established by Haskell, Hrushovski, and Macpherson in [15].

The purpose of this section is to give a very rough sketch of the latter proof.

(5.2) Canonical PL subsets of Berkovich spaces. Let X be a Hausdorff k -analytic space of dimension n .

(5.2.1) Let $x \in X$, let d be transcendence degree of $\widehat{\mathcal{H}}(x)$ over \widetilde{k} , and let r be the rational rank of the abelian group $|\mathcal{H}(x)^\times|/|k^\times|$. One always has the inequality $d + r \leq n$; the point x will be said to be *Abhyankar* if $d + r = n$.

(5.2.2) Let S be a locally closed subset of X . We will say that S is a *skeleton* if it consists of Abhyankar points and if it admits a PL structure satisfying the following conditions.

- For every analytic domain Y of X and every invertible function f on Y , the intersection $Y \cap S$ is a PL subspace of S , and the restriction of $|f|$ to $Y \cap S$ is PL.
- There exists a polyhedral atlas \mathfrak{A} on S having the following property: for every $P \in \mathfrak{A}$, the class of polyhedral structures defining the chart P admits a representative of the form $(|f_1|, \dots, |f_m|)|_P$ where the f_i ’s are invertible functions on an analytic domain Y of X which contains P .

If such a structure exists on S , it is easily seen to be unique. The archetypal example of a skeleton is S_n (5.1.3).

(5.2.3) A simple criterion for being a skeleton. The conditions for being a skeleton may seem slightly complicated to check. But in fact, in practice this is not so difficult. Indeed, let us assume that X is affinoid and irreducible, of dimension n , and let P be a compact subset of X consisting of Abhyankar points. The latter property ensures that every point of P is Zariski-generic. In particular, any nonzero analytic function on X is invertible on P .

Now assume that for every finite family (f_1, \dots, f_m) of nonzero analytic functions on X , the image $(|f_1|, \dots, |f_m|)(P)$ is a polyhedron of $(\mathbb{R}_+^\times)^m$, and that there exists such a family (f_1, \dots, f_m) with $(|f_1|, \dots, |f_m|)|_P$ injective. Under this assumption, P is a skeleton [10, Lemme 4.4].

The two main ingredients of the proofs are the Gerritzen–Grauert theorem (that describes the affinoid domains of X , see, for instance, [9, Lemme 2.4]), and the density of rational functions on X inside the ring of analytic functions of any rational affinoid domain of X .

(5.2.4) We can now rephrase Theorem (5.1.4) as follows.

Theorem 5.2.5 ([10, Theorem 5.1]). *If f is any morphism from X to $\mathbb{G}_m^{n,\text{an}}$, then $f^{-1}(S_n)$ is a skeleton, and $f^{-1}(S_n) \rightarrow S_n$ is a piecewise immersion.*

(5.3) The general strategy: algebraization and choice of a faithful tropicalization.

(5.3.1) **Algebraization.** In order to prove Theorem 5.2.5 we first algebraize the situation by standard arguments which we now outline (for more details, see [10, Sects. 5.4 and 5.5]).

First of all, one extends the scalars to the perfect closure of k , and then replaces X with its underlying reduced space; we can do it because the expected assertions are insensitive to radical extensions and nilpotents. Now X is generically quasi-smooth (quasi-smoothness is the Berkovich version of rig-smoothness, see [13, Sect. 4.2] et sq.). Every point of $f^{-1}(S_n)$ is Abhyankar (because so are the points of S_n and because $\dim X = n$), hence is quasi-smooth; one can thus shrink X so that it is itself quasi-smooth. Krasner’s lemma then ensures that X is G-locally algebraizable (see, for instance, [10, (0.21)]) which eventually allows, because being a skeleton is easily seen to be a G-local property, to reduce to the case where X is a connected, irreducible rational affinoid domain of \mathcal{X}^{an} for \mathcal{X} an irreducible, normal (and even smooth) algebraic variety over k .

By density arguments, we also can assume that f is induced by a dominant, generically finite algebraic map from \mathcal{X} to \mathbb{G}_m^n (which we still denote by f). Now let $x \in f^{-1}(S_n)$. It is Abhyankar hence Zariski-generic. By openness of finite, flat morphisms, it therefore admits a connected affinoid neighborhood U_x in \mathcal{X}^{an} such that f induces a finite and flat map from U_x to an affinoid domain V_x of $\mathbb{G}_m^{n,\text{an}}$ with $V_x \cap S_n$ being a nonempty n -dimensional simplex. Now there are finitely many such U_x ’s covering $X \cap f^{-1}(S_n)$. It is then sufficient to prove that for every $x \in S_n$ the intersection $U_x \cap f^{-1}(S_n)$ is a skeleton and that $U_x \cap f^{-1}(S_n) \rightarrow S_n$ is a piecewise immersion. Therefore we may assume that f induces a finite and flat map from X to an affinoid domain V of $\mathbb{G}_m^{n,\text{an}}$ with $V \cap S_n$ being a nonempty n -dimensional simplex.

The key result is now the following theorem.

(5.3.2) **Finite separation theorem.** *There exist finitely many nonzero rational functions g_1, \dots, g_r on \mathcal{X} whose norms separate the pre-images of x under f for every $x \in S_n$.*

We postpone the outline of the proof of Theorem (5.3.2) to paragraph (5.4); we are first going to explain how one can use it to prove Theorem (5.1.4), or more precisely its rephrasing written down at (5.2.4).

(5.3.3) **Pre-image of S_n and tropical dimension.** Let f_1, \dots, f_n be the invertible functions on \mathcal{X} that define f . For every compact analytic domain Y of \mathcal{X}^{an} , the subset $(|f_1|, \dots, |f_n|)(Y)$ is a polyhedron (cf. [6, Corollary 6.2.2]; see also [10, Theorem 3.2]).

If $x \in \mathcal{X}^{\text{an}}$, the *tropical dimension* of f at x is the infimum of the dimensions of the polyhedra $(|f_1|, \dots, |f_n|)(Y)$ for Y going through the set of compact analytic

neighborhoods of x . One can then characterize $f^{-1}(S_n)$ as the subset of \mathcal{X}^{an} consisting of points at which the tropical dimension of f is exactly n (this follows from [10, Theorem 3.4]).

(5.3.4) The compact subset $f^{-1}(S_n) \cap X$ of \mathcal{X}^{an} is a skeleton, and

$$U_x \cap f^{-1}(S_n) \rightarrow S_n$$

is a piecewise immersion. Since every point of $f^{-1}(S_n)$ is Zariski-generic, the functions g_i 's are invertible on $f^{-1}(S_n)$; hence we can shrink X so that the g_i 's are invertible on it. Let h_1, \dots, h_m be arbitrary nonzero analytic functions on X .

We set for short $|f| = (|f_1|, \dots, |f_n|): X \rightarrow (\mathbb{R}_+^{\times})^n$ and define $|g|$ and $|h|$ analogously. Let π be the map $(|f|, |g|, |h|)$ from X to $(\mathbb{R}_+^{\times})^{n+m+r}$; the image $\pi(X)$ is a polyhedron.

The restriction of $|f|$ to S_n is injective by the very definition of S_n , and $|g|$ separates the pre-images of every point of S_n on \mathcal{X}^{an} ; therefore, the restriction of $(|f|, |g|)$ to $f^{-1}(S_n)$ is injective; hence $\pi|_{f^{-1}(S_n) \cap X}$ is injective.

Let us choose a triangulation of the compact polyhedron $\pi(X)$ by convex compact polyhedra, and let P be the union of the n -dimensional closed cells Q such that the following holds: *the restriction of $(|f_1|, \dots, |f_n|)$ to Q is injective*. Using the characterization of $f^{-1}(S_n)$ through tropical dimension (5.3.3), the openness of finite, flat morphism and the fact that $V \cap S_n$ is nonempty of pure dimension n , one proves that $\pi(f^{-1}(S_n) \cap X) = P$ [10, 5.5.1]; in particular, this is a polyhedron and in view of (5.2.3) this implies that $f^{-1}(S_n) \cap X$ is a skeleton. Moreover modulo the isomorphism $(f^{-1}(S_n) \cap X) \simeq P$ induced by π , the function $|f|$ is nothing but the projection to the first n variables; it is PL and injective on each cell of P , which implies that $f^{-1}(S_n) \cap X \rightarrow S_n$ is a piecewise immersion. This ends the proof of Theorem 5.2.5.

(5.4) About the proof of the finite separation Theorem (5.3.2).

(5.4.1) This theorem asserts that there exist finitely many elements g_1, \dots, g_r of the function field $k(\mathcal{X}^{\circ})$ which separate the extensions of every real-valued Gauß norm on $k(T_1, \dots, T_n)$. In fact, we establish the more general, purely valuation-theoretic following theorem [10, Theorem 2.8]. Let k be an arbitrary valued field, let n be an integer, and let L be a finite extension of $k(T_1, \dots, T_n)$. There exists a finite subset E of L such that the following hold: *for every ordered abelian group G containing $|k|$ and every element $r = (r_1, \dots, r_n)$ of G^n , the elements of E separate the extensions of the G -valued valuation η_r of $k(T_1, \dots, T_n)$ to L .*

(5.4.2) A particular case. In order to illustrate the general idea the proof is based upon, let us explain what is going on in a simple particular case, namely, if k is algebraically closed, and if $n = 1$; we write T instead of T_1 . The field L can then be written $k(C)$ for a suitable projective, smooth, irreducible curve C over k , equipped with a finite morphism $C \rightarrow \mathbb{P}_k^1$ inducing the extension $k(T) \hookrightarrow L$. The map $C \rightarrow \mathbb{P}_k^1$ induces a k -definable natural transformation $\widehat{f} : \widehat{C} \rightarrow \widehat{\mathbb{P}_k^1}$ at the level of stable completions; it follows from Theorem 3.10 that \widehat{C} and $\widehat{\mathbb{P}_k^1}$ are k -definable (for $\widehat{\mathbb{P}_k^1}$ this can also be seen directly, cf. (3.8)).

The map $r \mapsto \eta_r$ defines a k -definable natural embedding $\Gamma \hookrightarrow \widehat{\mathbb{P}}_k^1$; let Δ be the pre-image of Γ under \widehat{f} . This is a sub-functor of \widehat{C} which is k -definable by its very definition, and the fibers of $\Delta \rightarrow \Gamma$ are finite.

Let F be an algebraically closed valued extension of k , and let $r \in |F|$. The pre-image D of $\eta_{r,F}$ inside $\Delta(F)$ is by construction definable over the set of parameters $k \cup \{r\}$, and is finite. By a result proven in [15], this implies that every element of D is *individually* $(k \cup \{r\})$ -definable.⁵ In other words, for every $d \in D$, there exists a k -definable natural transformation $\sigma : \Gamma \rightarrow \Delta$ such that $d = \sigma(\eta_{r,F})$.

As this holds for arbitrary F and r , the celebrated compactness theorem of model theory ensures the existence of finitely many sections $\sigma_1, \dots, \sigma_m$ of $\Delta \rightarrow \Gamma$ such that $\Delta = \bigcup \sigma_i(\Gamma)$. Every sub-functor $\sigma_i(\Gamma)$ is k -definable, and k -definably isomorphic to Γ through σ_i . One then easily builds, starting from the equality $\Delta = \bigcup \sigma_i(\Gamma)$, a k -definable embedding $\Delta \hookrightarrow \Gamma^N$ (for some big enough N —one simply has to be able to realize the “coincidence diagram” of the σ_i ’s inside Γ^N).

Now Hrushovski and Loeser have proven that S is a k -definable sub-functor of \widehat{C} such that there exists a k -definable embedding $S \hookrightarrow \Gamma^N$ then there exists such an isomorphism *induced by the norms of finitely many k -rational functions on \widehat{C}* [16, Proposition 6.2.7]; this comes from the explicit description of \widehat{C} as a k -definable functor, and from general results established in [15]. Applying this to Δ , we get the existence of finitely many rational functions on C whose norms separate points of Δ ; in particular, those functions separate for every algebraically closed valued extension F of k and every $r \in |F^\times|$ the pre-images of $\eta_{r,F}$ in $\widehat{C}(F)$. They *a fortiori* separate the extensions of the valuation $\eta_{r,k}$ to the field L because by elementary valuation-theoretic arguments, such an extension is always the restriction of a valuation on $F(C_F) = L \otimes_{k(T)} F(T)$ inducing $\eta_{r,F}$ on $F(T)$.

(5.4.3) The general case. The proof goes by induction on n . The crucial step is of course the one that consists in going from $n - 1$ to n , and it roughly consists of a relative version of what we have done above.

(5.5) Some complements. We still denote by k a complete, non-Archimedean field and by X an n -dimensional k -analytic space.

(5.5.1) Finite union of pre-images of the skeleton. Let f_1, \dots, f_m be morphisms from X to $\mathbb{G}_m^{n,\text{an}}$.

Theorem 5.2.5 states that $f_i^{-1}(S_n)$ is a skeleton for every i . Theorem 5.1 of [10] in fact also states that the union $\bigcup_i f_i^{-1}(S_n)$ is still a skeleton; this essentially means that for every (i, j) , the intersection of the skeletons $f_i^{-1}(S_n)$ and $f_j^{-1}(S_n)$ is PL in both of them. The proof consists in exhibiting in a rather explicit way an integer N ,

⁵It is crucial for this result that the additional parameter r belongs to the value group of F , and not to F itself. For instance, let a be an element of F which is transcendental over k , assume that $\text{char. } k \neq 2$ and let E be the two-element set of square roots of a . Then E is globally definable over $k \cup \{a\}$, but this is not the case of any of those two square roots: one cannot distinguish between them by a formula involving only a and elements of k .

a “Shilov section with non-constant radius” $\sigma : X \rightarrow \mathbb{A}_X^N$, and a skeleton $P \subset \mathbb{A}_X^N$ (described as the pre-image of S_{N+n} under a suitable map) such that for every i , the section σ identifies $f_i^{-1}(S_n)$ with a PL subspace of P [10, 5.6.2].

(5.5.2) Stabilization after a finite, separable extension. For every complete extension F of k , let $\Sigma_F \subset X_F$ be the skeleton $\bigcup_i f_{i,F}^{-1}(S_{n,F})$, where $S_{n,F}$ is the skeleton of $\mathbb{G}_{m,F}^{n,\text{an}}$.

If $F \hookrightarrow L$, there is a natural surjection $\Sigma_L \rightarrow \Sigma_F$, which is a piecewise immersion of PL spaces. Theorem 5.1 of [10] in fact also states that there exists a finite separable extension F of k such that $\Sigma_L \rightarrow \Sigma_F$ is a homeomorphism for every complete extension L of F . To see it, one essentially refines the proof of the “separation theorem” on Gauß norms that is sketched at (5.4) et sq. to show that after making a finite, separable extension of the ground field, one can exhibit a finite family of functions separating *universally* (that is, after any extension of the ground field) the extensions of Gauß norms; this is part of aforementioned Theorem 2.8 of loc. cit.

Acknowledgements I’m grateful to the referees for their careful reading, and for many helpful comments and suggestions.

During this work the author was partially supported by ValCoMo (ANR-13-BS01-0006), and by a Rosi and Max Varon visiting professorship at Weizmann Institute.

References

1. Abbes, A.: Ramification of local fields with imperfect residue fields. *Am. J. Math.* **124**(5), 879–920 (2002)
2. Berkovich, V.: Spectral theory and analytic geometry over non-Archimedean fields. In: *Mathematical Surveys and Monographs*, vol. 33. AMS, Providence, RI (1990)
3. Berkovich, V.: Étale cohomology for non-Archimedean analytic spaces. *Inst. Hautes Etudes Sci. Publ. Math.* **78**, 5–161 (1993)
4. Berkovich, V.: Vanishing cycles for formal schemes. *Invent. Math.* **115**, 539–571 (1994)
5. Berkovich, V.: Smooth p-adic analytic spaces are locally contractible. *Invent. Math.* **134**, 1–84 (1999)
6. Berkovich, V.: Smooth p-adic analytic spaces are locally contractible. II. In: *Geometric aspects of Dwork theory*, vols. I, II, pp. 293–370. Walter de Gruyter GmbH & Co. KG, Berlin (2004)
7. Chatzidakis, Z.: Théorie des modèles des corps valués. Master’s course. Available at <http://www.logique.jussieu.fr/~zoe>
8. Ducros, A.: Image réciproque du squelette par un morphisme entre espaces de Berkovich de même dimension. *Bull. Soc. Math. Fr.* **13**(4), 483–506 (2003)
9. Ducros, A.: Parties semi-algébriques d’une variété algébrique p-adique. *Manuscripta Math.* **111**(4), 513–528 (2003)
10. Ducros, A.: Espaces de Berkovich, polytopes, squelettes et théorie des modèles. *Confluentes Math.* **4**(4) (2012). Available at <http://math.univ-lyon1.fr/~confluentes/CM/04/0404/Ducros.pdf>
11. Ducros, A.: Espaces de Berkovich, polytopes, squelettes et théorie des modèles: erratum. *Confluentes Math.* **5**(2), 43–44 (2013)

12. Ducros, A.: Les espaces de Berkovich sont modérés (d'après Ehud Hrushovski et François Loeser) [French]. [Berkovich spaces are tame (following Ehud Hrushovski and François Loeser)]. In: Séminaire Bourbaki, vol. 2011/2012, Exposés 1043-1058. Astérisque, vol. 352 (2013), Exp. No. 1056, x, pp. 459–507 [ISBN: 978-2-85629-371-3] [Reviewer: Antoine Chambert-Loir, 14G22 (03C60)]
13. Ducros, A.: Families of Berkovich spaces. Preprint, arXiv:1107.4259
14. Haskell, D., Hrushovski, E., Macpherson, D.: Definable sets in algebraically closed valued fields: elimination of imaginaries. *J. Reine Angew. Math.* **597**, 175–236 (2006)
15. Haskell, D., Hrushovski, E., Macpherson, D.: Stable domination and independence in algebraically closed valued fields. In: *Lecture Notes in Logic*, vol. 30. Association for Symbolic Logic/Cambridge University Press, Chicago, IL/Cambridge (2008)
16. Hrushovski, E., Loeser, F.: Non-Archimedean Tame Topology and Stably Dominated Types. *Annals of Mathematics Studies*, vol. 192, vii+216 pp. Princeton University Press, Princeton (2016) [MR3445772]
17. Poineau, J.: Un résultat de connexité pour les variétés analytiques p -adiques: privilège et noethérianité. *Compos. Math.* **144**(1), 107–133 (2008)
18. Poineau, J.: Polynomial approximation of Berkovich spaces and definable types. *Math. Ann.* **358**(3–4), 949–970 (2014)
19. Prestel, A.: *Einführung in die Mathematische Logik und Modelltheorie*. Vieweg, Braunschweig (1986)
20. Raynaud, M.: Géométrie analytique rigide d'après Tate, Kiehl,... In: *Table Ronde d'Analyse non archimédienne* (Paris, 1972). *Bulletin de la Société Mathématique de France*, Mem. No. 39–40, pp. 319–327. Society of Mathematics, France (1974)
21. Robinson, A.: Complete theories. In: *Studies in Logic and the Foundations of Mathematics*, vol. 13. North Holland, Amsterdam (1956)
22. Tate, J.: Rigid analytic spaces. *Invent. Math.* **12**, 257–289 (1971)
23. Weispfenning, V.: Quantifier elimination and decision procedures for valued fields. In: *Models and Sets*, Aachen, 1983. *Lecture Notes in Mathematics*, vol. 1103, pp. 419–472. Springer, Berlin 1984

Excluded Homeomorphism Types for Dual Complexes of Surfaces

Dustin Cartwright

Abstract We study an obstruction to prescribing the dual complex of a strict semistable degeneration of an algebraic surface. In particular, we show that if Δ is a complex homeomorphic to a 2-dimensional manifold with negative Euler characteristic, then Δ is not the dual complex of any semistable degeneration. In fact, our theorem is somewhat more general and applies to a certain class of complexes homotopy equivalent to such a manifold. Our obstruction is provided by the theory of tropical complexes.

Keywords Semistable • Degeneration • Dual complex

1 Introduction

The dual complex of a semistable degeneration is a combinatorial encoding of the combinatorics of the components of the special fiber. In recent years, it has been studied because of connections to tropical geometry [11], non-Archimedean analytic geometry [2], and birational geometry [3, 6]. In this paper, we study obstructions to realizing arbitrary complexes as dual complexes of degenerations of surfaces.

We let R be any rank 1 valuation ring with algebraically closed residue field and we will consider a *degeneration* over R to be a flat, proper scheme \mathfrak{X} over $\text{Spec } R$ which is strictly semistable in the sense of [9, Sect. 3]. Specifically, we require that \mathfrak{X} is covered by open sets which admit étale morphisms over R to $\text{Spec } R[x_0, \dots, x_n]/(x_0 \cdots x_m - \pi)$ for some $m \leq n$ and some non-zero π in the maximal ideal of R . The dual complex of \mathfrak{X} is a Δ -complex with one vertex for each irreducible component of the special fiber and higher-dimensional simplices for each connected component where irreducible components intersect.

Since semistability implies that the special fiber is normal crossing, the dimension of the dual complex is at most the relative dimension of the family \mathfrak{X} . In dimension 1, any graph is the dual complex of some degeneration of curves

D. Cartwright (✉)

Department of Mathematics, University of Tennessee, 227 Ayres Hall, Knoxville, TN 37996-1320, USA

e-mail: cartwright@utk.edu

[1, Corollary B.3]. However, in this paper, we show that the analogous statement is not true in dimension 2.

Theorem 1.1. *There is no strict semistable degeneration over a rank 1 valuation ring R , whose general fiber is a smooth, geometrically irreducible surface and such that the dual complex of the special fiber is homeomorphic to a topological surface Σ with $\chi(\Sigma) < 0$.*

We conjecture that Theorem 1.1 can be strengthened to replace “homeomorphic” with “homotopy equivalent.” In fact, we can prove a strengthening in this direction which applies to Δ -complexes formed from a manifold with negative Euler characteristic by attaching additional simplices in a controlled way. First, what we call “fins” are allowed so long as they don’t change the homotopy type and where the gluing is along a subset that’s not too complicated. Second, arbitrary complexes may be attached to the manifold, so long as the gluing is along a finite set. These complexes are collectively the “ornaments” in the following definition.

Definition 1.2. We say that a 2-dimensional Δ -complex Δ is a *manifold with fins and ornaments* if there exists a subcomplex Σ , the *manifold*, subcomplexes F_1, \dots, F_n , the *fins*, and a subcomplex O , the *ornaments*, such that:

- (1) We have a decomposition $\Delta = \Sigma \cup F_1 \cup \dots \cup F_n \cup O$.
- (2) Σ is homeomorphic to a connected 2-dimensional topological manifold.
- (3) For any i , F_i is contractible and $F_i \cap \Sigma$ is a path.
- (4) For any $i > j$, $F_j \cap F_i$ is a subset of the endpoints of the path $F_i \cap \Sigma$.
- (5) The intersection $O \cap (\Sigma \cup F_1 \cup \dots \cup F_n)$ is a finite set of points.

If the manifold Σ has negative Euler characteristic, we call Δ a *hyperbolic manifold with fins and ornaments* and if the subcomplex O is empty, then we call Δ a *manifold with fins*.

Theorem 1.3. *If Δ is a hyperbolic manifold with fins and ornaments, then there is no degeneration with dual complex Δ .*

The obstruction used in the proof of Theorem 1.3 is in lifting the dual complex to a tropical complex. A tropical complex is a Δ -complex, together with the intersection numbers of the 1-dimensional strata inside the 2-dimensional strata, which are called the structure constants of the tropical complex [5].

Theorem 1.4. *If Δ is a hyperbolic manifold with fins and ornaments, then there is no 2-dimensional tropical complex with Δ as its underlying topological space.*

Note that when we construct the tropical complex from a degeneration, we do not incorporate valuations from the defining equations, in contrast with both the suggestion from the introduction of [5] and the construction in [9, Example 3.10], where edges of the dual complex have lengths coming from the value group of R . We believe that for many applications, such metric information will be essential, but in this paper, we intentionally ignore it for the purpose of being able to use the results from [4], where the edges implicitly all have length 1. Philosophically, we think of our approach as taking a deformation in the category of “metric tropical

complexes” from the complex which encodes the valuations to a tropical complex with all edges of length 1, where we can apply Theorem 1.4.

There are two characteristics of the special fiber of a degeneration which are incorporated into the axioms of a tropical complex. The first is that the special fiber is principal which gives a relationship among the intersection numbers with a fixed curve. The second is that Hodge index theorem, which restricts the possible intersection matrices of a fixed surface in the special fiber.

Both axioms of a tropical complex are necessary in the proof of the obstruction. Without the condition coming from the Hodge index theorem, the object would only be a weak tropical complex, and any Δ -complex lifts to a weak tropical complex. For example, if Δ is homeomorphic to a topological manifold, then choosing all structure constants equal to 1 gives a weak tropical complex, but this will not be a tropical complex if $\chi(\Delta) < 0$, as explained in Example 2.2.

On the other hand, Kollár has shown that any finite n -dimensional simplicial complex is realizable as the dual complex of a simple normal crossing divisor [13, Theorem 1], but such a divisor would not give a tropical complex because the divisor is not necessarily principal. However, when connected, such a divisor can be realized as the exceptional locus of the resolution of a normal, isolated singularity [13, Theorem 2]. Thus, we see Theorem 1.3 as an example of how the global geometry of a smooth algebraic variety is more restricted than the local geometry of a singularity, in line with [12].

We also note that unlike the cases in Theorem 1.1, topological surfaces with non-negative Euler characteristic are all possible as homeomorphism types of degenerations. In particular, the 2-sphere, the real projective plane, the torus, and the Klein bottle appear as degenerations of K3 surfaces, Enriques surfaces, Abelian surfaces, and bielliptic surfaces, respectively. In fact, a partial converse is possible in that the dual complexes of such degenerations have been classified by results of Kulikov, Persson, Pinkham, and Morrison [14–17]. Note that topological surfaces of non-negative Euler characteristic all arose from degenerations of varieties of Kodaira dimension 0. However, these classification results would already suffice to prove Theorem 1.1 if we assumed that the general fiber had Kodaira dimension 0.

2 Tropical Complexes and Tropical Surfaces

We begin by recalling the definition of tropical complexes, as introduced in [5] and their properties, as studied in [4]. Unlike those papers, we additionally assume that the underlying Δ -complex is regular, meaning that the faces of any fixed simplex are distinct. All of our combinatorial results also hold without a regularity assumption, but the dual complex of a strictly semistable degeneration is always regular, so that case is sufficient for our applications. In addition, we will only work with 2-dimensional tropical complexes, which we will call tropical surfaces. We will refer to the simplices of dimensions 0, 1, and 2 in a 2-dimensional Δ -complex as its *vertices*, *edges*, and *facets*, respectively.

Definition 2.1. A *weak tropical surface* is a finite, connected, regular Δ -complex whose cells have dimension at most 2, together with integers $\alpha(v, e)$ for every endpoint v of an edge e , such that for each edge e , we have an equality:

$$\alpha(v, e) + \alpha(w, e) = \deg(e), \quad (1)$$

where v and w are the endpoints of e and $\deg(e)$ is the number of 2-dimensional facets containing e .

At a vertex v of a weak tropical surface Δ , the *local intersection matrix* M_v is a symmetric matrix whose rows and columns are indexed by edges containing v and such that the entry corresponding to edges e and e' is

$$(M_v)_{e,e'} = \begin{cases} \#\{\text{facets containing both } e \text{ and } e'\} & \text{if } e \neq e' \\ -\alpha(w, e) & \text{if } e = e', \end{cases} \quad (2)$$

where w is the endpoint of e other than v . A *tropical surface* is a weak tropical surface Δ such that for every vertex v of Δ , M_v has exactly one positive eigenvalue.

Example 2.2. Let Δ be a triangulated manifold, meaning a regular Δ -complex which is homeomorphic to a 2-dimensional, connected topological manifold. Then, every edge e is contained in two facets, so a symmetric choice for the structure constants satisfying the constraint (1) on e is to set $\alpha(v, e) = \alpha(w, e) = 1$ for both endpoints v and w . This gives us a weak tropical surface.

At any vertex v , the link is a cycle and so, for example, if this cycle has length 5, the local intersection matrix is

$$M_v = \begin{pmatrix} -1 & 1 & 0 & 0 & 1 \\ 1 & -1 & 1 & 0 & 0 \\ 0 & 1 & -1 & 1 & 0 \\ 0 & 0 & 1 & -1 & 1 \\ 1 & 0 & 0 & 1 & -1 \end{pmatrix}$$

An $m \times m$ matrix of this form always has a positive eigenvalue of 1 for the eigenvector $(1, \dots, 1)^T$. However, if $m > 6$, then there exist at least two other positive eigenvalues. Therefore, Δ is a tropical complex if and only if each vertex is contained in at most six edges. A simple counting argument shows that if each vertex is contained in at most six edges, then the Euler characteristic $\chi(\Delta)$ is non-negative. Thus, we've verified the special case of Theorem 1.4 that there is no tropical surface with all structure constants equal to 1 for which the underlying topological space is a hyperbolic manifold.

In this paper, the main purpose of the structure constants $\alpha(v, e)$ is to define a sheaf of linear functions on a weak tropical surface Δ . These linear functions are the tropical analogue of non-vanishing regular functions in algebraic geometry.

Definition 2.3 (Construction 4.2 and Definition 4.3 in [5]). Let Δ be a weak tropical surface and e an edge of Δ . We define N_e to be the simplicial complex consisting of e with a 2-dimensional simplex attached for each facet of Δ containing e . Then, there exists a natural map $\pi_e: N_e \rightarrow \Delta$ which is an open inclusion on N_e^o , which we define to be the union of interior of e and of the interiors of the facets of N_e .

We also define a map $\phi_e: N_e \rightarrow \mathbb{R}^{d+2}/\mathbb{R}$, where d is the number of facets of Δ containing e and the quotient is by the line in \mathbb{R}^{d+2} generated by the vector $(1, \dots, 1, -\alpha(v, e), -\alpha(w, e))$. The map ϕ_e sends each vertex of N_e to the image \mathbf{e}_i of the i th basis vector of \mathbb{R}^{d+2} , with v and w going to \mathbf{e}_{d+1} and \mathbf{e}_{d+2} , respectively. We then extend ϕ_e linearly to all of N_e . The vectors $\mathbf{e}_1, \dots, \mathbf{e}_{d+2}$ generate a lattice inside $\mathbb{R}^{d+2}/\mathbb{R}$ and we say that a function $\ell: \mathbb{R}^{d+2}/\mathbb{R} \rightarrow \mathbb{R}$ has linear slopes if $\ell(\mathbf{e}_i) - \ell(\mathbf{e}_j)$ is an integer for any two vectors \mathbf{e}_i and \mathbf{e}_j .

Finally, we say that a continuous function ϕ on an open subset $U \subset \Delta$ is *linear* if the following two conditions hold. First, if we identify the interior of any facet f meeting U with a unimodular simplex in \mathbb{R}^2 , then ϕ is an affine linear function with integral slopes on the interior of f . Second, on any edge e meeting U , $\phi|_{U \cap N_e} = \ell \circ \phi_e$ for some linear function $\ell: \mathbb{R}^{d+2}/\mathbb{R} \rightarrow \mathbb{R}$ with integral slopes. We write \mathcal{A} for the sheaf of linear functions on Δ .

Note that the image of the map ϕ_e in Definition 2.3 lies in the affine linear space consisting of the linear combinations $\sum c_i \mathbf{e}_i$ such that $c_1 + \dots + c_{d+2} = 1$, which is a proper subset of $\mathbb{R}^{d+2}/\mathbb{R}$ because of relation (1) in the definition of a weak tropical surface. Roughly speaking, the quotient in Definition 2.3 means that linearity on a weak tropical surface imposes one condition beyond linearity on each facet of Δ . More precisely, we have the following:

Lemma 2.4. *Let e be an edge of a tropical surface Δ and let N_e^o be the union of the interiors of e and of the facets containing e , as in Definition 2.3. If ϕ is a linear function on N_e^o , and is constant on all but one of these facets, then ϕ is constant.*

Proof. If ϕ_e is as in Definition 2.3, then ϕ factors as $\ell \circ \phi_e$ for some affine linear function ℓ . We also let d and \mathbf{e} be as in Definition 2.3, and we can assume that the facets containing e are numbered so that ϕ is constant on all but the first one. Then, we have the linear relation

$$\mathbf{e}_1 = -\mathbf{e}_2 - \dots - \mathbf{e}_d + \alpha(v, e)\mathbf{e}_{d+1} + \alpha(w, e)\mathbf{e}_{d+2},$$

where the sum of the coefficients is 1 by the relation (1), and so $\ell(\mathbf{e}_2) = \dots = \ell(\mathbf{e}_{d+2})$ implies that $\ell(\mathbf{e}_1) = \ell(\mathbf{e}_2) = \dots = \ell(\mathbf{e}_{d+2})$. Therefore, ℓ is constant on $\phi_e(N_e^o)$, which is what we wanted to show. \square

Example 2.5. As in Example 2.2, suppose we have a triangulated manifold with $\alpha(v, e) = 1$ for all endpoints v of all edges e . Then, in Definition 2.3, the interior of e and of the two facets containing e is identified with the interior of a unit square in \mathbb{R}^2 , with the square triangulated along a diagonal. Linear functions on this neighborhood of e are equivalent to affine linear functions on the square. Thus, in this case, Lemma 2.4 amounts to the observation that affine linear functions on \mathbb{R}^2 are determined by their restriction to any open set.

Any constant function on a weak tropical surface is linear, so the sheaf of locally constant \mathbb{R} -valued functions is a subsheaf of \mathcal{A} . If we denote the quotient sheaf \mathcal{A}/\mathbb{R} by \mathcal{D} , then we have a long exact sequence in sheaf cohomology [4, Sect. 3]:

$$0 \rightarrow H^0(\Delta, \mathbb{R}) \rightarrow H^0(\Delta, \mathcal{A}) \rightarrow H^0(\Delta, \mathcal{D}) \rightarrow H^1(\Delta, \mathbb{R}) \rightarrow \dots \quad (3)$$

One of the main results from [4] is the following:

Theorem 2.6 (Theorem 4.7 in [4]). *If Δ is a tropical surface which is locally connected through codimension 1, the \mathbb{R} -span of the image of the morphism $H^0(\Delta, \mathcal{D}) \rightarrow H^1(\Delta, \mathbb{R})$ has codimension at most 1 in $H^1(\Delta, \mathbb{R})$.*

In Theorem 2.6, *locally connected through codimension 1* means that the link of each vertex is connected.

If linear functions on weak tropical surfaces are taken to be analogous to non-vanishing regular functions, then the analogues of rational functions on an algebraic variety come from relaxing the linearity condition in codimension 1. More precisely, suppose that d_1, \dots, d_m are closed line segments, each in one facet of a weak tropical surface Δ , such that $d_i \cap d_j$ is finite for distinct i and j . Then we say a function ϕ on an open subset $U \subset \Delta$ is a *piecewise linear function whose divisor is supported in $d_1 \cup \dots \cup d_m$* if ϕ is continuous and ϕ is linear on $U \setminus (d_1 \cup \dots \cup d_m)$. Although we will not need it in this paper, we can justify our terminology with:

Proposition 2.7 (Proposition 4.5 in [5]). *Let d_1, \dots, d_m be line segments in a weak tropical surface Δ as above. If U is an open set meeting all of the d_i , then there exists a homomorphism from the group of piecewise linear functions ϕ whose divisor is supported on $d_1 \cup \dots \cup d_m$ under addition to formal sums of the d_i , known as the divisor of ϕ . The maximal open subset of U on which such a function ϕ is linear is the complement of those d_i with non-zero coefficient in the divisor of ϕ .*

Finally, we have an analogue of the maximum modulus principle from complex analysis. The result in [4] also applies to piecewise linear functions whose divisor has non-negative coefficients, but we'll only need it for linear functions, i.e., when the divisor is trivial.

Proposition 2.8 (Proposition 2.11 in [4]). *Let Δ be a tropical surface which is connected through codimension 1. If ϕ is a linear function on a connected open set U which achieves its maximum on U , then ϕ is constant.*

3 Degenerations

In this section, we construct weak tropical surfaces from strictly semistable degenerations. The case of regular semistable degenerations over discrete valuation rings was treated in [5, Sect. 2], but here we want to work over possibly non-discrete

valuation rings. The data of a weak tropical surface only depends on the special fiber as a simple normal crossing scheme over the residue field, and not on the valuation ring, so we can use the same construction as [5, Sect. 2], which we now recall.

We let \mathfrak{X} be a strictly semistable degeneration over R and, as stated in the introduction, the dual complex Δ has one k -dimensional simplex for each stratum of dimension $2 - k$ in the special fiber of \mathfrak{X} . Thus, if e is an edge with endpoints v and w , then we let C_e and C_w denote the curve and surface corresponding to e and w , respectively. We set $\alpha(v, e) = -C_e^2$, where C_e^2 denotes the self-intersection number of C_e in C_w . Even for regular strictly semistable degenerations over discrete valuation rings, this data may only give a weak tropical surface without an additional technical condition of robustness in dimension 2 [5, Proposition 2.7]. Over non-discrete valuation rings, we also get weak tropical surfaces.

Proposition 3.1. *The special fiber of any degeneration \mathfrak{X} yields a weak tropical surface Δ such that the local intersection matrix M_v has at most one positive eigenvalue for each vertex v of Δ .*

Proof. For Δ to be a weak tropical surface, we need to check that for any edge e with endpoints v and w , we have the equality (1):

$$\alpha(v, e) + \alpha(w, e) = \deg e. \tag{4}$$

Let C_e be the curve corresponding to e in the special fiber of \mathfrak{X} , and by our semistability condition, on a Zariski open neighborhood meeting C_e , there is an étale map to $\text{Spec } R[x, y, z]/\langle xy - \pi \rangle$ for some element π in the maximal ideal of R . Then, the principal Cartier divisor defined by π can be written, at least in a formal open neighborhood of C_e , as the union of Cartier divisors, each of which is supported on an irreducible component of the special fiber of \mathfrak{X} . For example, in the above chart, the functions x and y pull back to give defining equations for each of the components containing C_e .

Thus, using linearity of the intersection product [8, Proposition 5.9(b)], we can split up the intersection of the principal divisor defined by π with the curve C_e into terms coming from the components of the special fiber of \mathfrak{X} . For components of the special fiber which don't contain C_e , if we pull back to C_e we get a Cartier divisor equal to the points of intersection, with multiplicities equal to 1. Thus, the degree of the intersection of such a Cartier divisor with C_e is equal to the number of points of intersection by the projection formula [8, Proposition 5.9(c)]. The total degree for all components which don't contain C_e gives $\deg e$, which is the right-hand side of (4).

Now consider the two components C_v and C_w containing C_e . If we pull back the Cartier divisor supported on C_v to C_w , then we get the divisor C_e on C_w . The self-intersection of C_e is $-\alpha(v, e)$ by the definition of the structure constants. Thus, using the projection formula again, the components containing C_e contribute a cycle of degree equal to $-\alpha(v, e) - \alpha(w, e)$, so the desired equality (4) follows because π obviously defines a principal divisor.

Finally, we claim that the local intersection matrix M_v records the intersection theory on the surface of the special fiber corresponding to v , restricted to curves of the special fiber. For the diagonal entries of the local intersection matrix (2), this follows immediately from our definition of the structure constants. The off-diagonal entries of the local intersection matrix count facets containing two edges e and e' , which are in bijection with the number of reduced points in the intersection of corresponding curves C_e and $C_{e'}$, and thus equal to the intersection number $C_e \cdot C_{e'}$. Therefore, by the Hodge index theorem, the local intersection matrix M_v can have at most one positive eigenvalue. \square

One approach to obtaining a tropical surface instead of a weak tropical surface is Proposition 2.10 in [5], which shows that for degenerations with projective components, robustness can be obtained by appropriate blow-ups. Rather than adapting this proposition to the case of non-discrete valuations, while also keeping track of the effect on the underlying topological space, it is more convenient to perform the modification combinatorially:

Lemma 3.2. *Let Δ be a weak tropical surface and suppose that for each vertex v of Δ , the local intersection matrix M_v has at most one positive eigenvalue. Then, there exists a tropical surface Δ' such that the underlying topological space of Δ' is formed by attaching a finite number of 2-simplices to edges of Δ .*

Proof. We suppose that v is a vertex of Δ such that M_v has no positive eigenvalues, i.e., it is negative semidefinite. Let e be an edge containing v and let w be the other endpoint of e . We attach an additional 2-simplex onto e and label the new vertex u' , with the new edges e'_v and e'_w . We use v' , w' , and e' to denote the representatives of v , w , and e in the new weak tropical surface Δ' . We assign the coefficients on Δ' to be the same as on Δ , except that

$$\begin{aligned} \alpha(w', e') &= \alpha(w, e) & \alpha(v', e') &= \alpha(v, e) + 1 \\ \alpha(w', e'_w) &= 0 & \alpha(u', e'_w) &= 1 \\ \alpha(v', e'_v) &= 2 & \alpha(u', e'_v) &= -1 \end{aligned}$$

Then we claim the local intersection matrices $M_{u'}$ and $M_{v'}$ each have exactly one positive eigenvalue and that the number of positive eigenvalues of $M_{w'}$ is the same as that of M_w . Once we show the claim, then we can repeat the above construction at each vertex v whose local intersection matrix M_v is negative semidefinite to get the desired tropical surface.

The first part of the claim is that the local intersection matrices

$$M_{u'} = \begin{pmatrix} 0 & 1 \\ 1 & -2 \end{pmatrix} \text{ and } M_{v'} = \begin{pmatrix} * \cdots & * & 0 \\ \vdots & \vdots & \vdots \\ * \cdots & -\alpha(w, e) & 1 \\ 0 \cdots & 1 & 1 \end{pmatrix}$$

have exactly one positive eigenvalue each, where $*$ indicates the parts of $M_{v'}$ that coincide with M_v . For $M_{u'}$, there is exactly one positive eigenvalue because it has determinant -1 . For $M_{v'}$, we use the change of coordinates:

$$NM_{v'}N^T = \begin{pmatrix} * \cdots & * & 0 \\ \vdots & \vdots & \vdots \\ * \cdots & -\alpha(w, e) - 1 & 0 \\ 0 \cdots & 0 & 1 \end{pmatrix} \text{ where } N = \begin{pmatrix} 1 \cdots 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 \cdots 1 & -1 \\ 0 \cdots 0 & 1 \end{pmatrix}$$

to get a block diagonal matrix whose upper left block is M_v minus a negative semidefinite diagonal matrix, and is thus negative semidefinite. Moreover, the lower right block of $NM_{v'}N^T$ is a single positive entry, so $NM_{v'}N^T$ has exactly one positive eigenvalue, as in the first part of the claim.

For the second part of our claim, we want to show that

$$M_{w'} = \begin{pmatrix} * \cdots & * & 0 \\ \vdots & \vdots & \vdots \\ * \cdots & -\alpha(v, e) - 1 & 1 \\ 0 \cdots & 1 & -1 \end{pmatrix}$$

has the same number of positive eigenvalues as M_w , where $*$ denotes parts of the matrix which coincide with M_w . We again use a change of coordinates to a block diagonal matrix:

$$PM_{w'}P^T = \begin{pmatrix} * \cdots & * & 0 \\ \vdots & \vdots & \vdots \\ * \cdots & -\alpha(v, e) & 0 \\ 0 \cdots & 0 & -1 \end{pmatrix} \text{ where } P = \begin{pmatrix} 1 \cdots 0 & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 \cdots 1 & 1 \\ 0 \cdots 0 & 1 \end{pmatrix}.$$

The upper left block of $PM_{w'}P^T$ agrees with M_w and the lower right block adds a single negative eigenvalue. Thus, $M_{w'}$ has the same number of positive eigenvalues as M_w , which completes the proof of the claim. \square

4 Proof of the Main Theorems

The crux of Theorem 1.4 and thus of Theorem 1.3 is the following lemma:

Lemma 4.1. *Let Δ be a tropical surface whose underlying Δ -complex is a manifold with fins and is connected through codimension 1. If s is a facet contained in*

the manifold subcomplex of Δ , and U_s denotes the interior of s , then the restriction map

$$H^0(\Delta, \mathcal{D}) \rightarrow H^0(U_s, \mathcal{D}) \cong \mathbb{Z}^2$$

is injective.

Proof. Note that the isomorphism $H^0(U_s, \mathcal{D}) \cong \mathbb{Z}^2$ holds because affine linear functions on U_s are equivalent to affine linear functions with integral slopes on a unimodular simplex in \mathbb{R}^2 by definition. Thus, \mathcal{A} restricted to U_s is isomorphic to the locally constant sheaf with values in $\mathbb{R} \times \mathbb{Z}^2$, and the quotient sheaf $\mathcal{D} = \mathcal{A}/\mathbb{R}$ is isomorphic to \mathbb{Z}^2 .

Now, we let Σ and F_1, \dots, F_n denote the manifold and fins of the simplicial complex, as in Definition 1.2. We suppose ω is a global section of \mathcal{D} such that the restriction of ω to U_s is trivial, and we want to show that ω is trivial. We start with $V = U_s$ and then we'll expand the open set V until it is all of Δ . At each step, V will either be disjoint from each fin F_i or contain F_i , except possibly the endpoints of $F_i \cap \Sigma$. In particular, the boundary of V will be contained in Σ .

First suppose that there exists an edge e in the boundary of V such that e is not contained in any of the fins. Let f denote the 2-simplex bordering e whose interior is in V and let f' denote the 2-simplex on the other side of e . Then, Definition 2.3 identifies the union of the interiors of f, f' , and e with an open subset of \mathbb{R}^2 , and the sections of \mathcal{A} are exactly affine linear functions with integral slope on this set. As above, \mathcal{A} and \mathcal{D} are therefore locally constant sheaves with values in $\mathbb{R} \times \mathbb{Z}^2$ and \mathbb{Z}^2 , respectively. Thus, we can expand V to include the interiors of e and f' , where ω is also zero.

Second, we assume that every edge in the boundary of V is contained in some $\Sigma \cap F_i$. Let i be the maximal index such that F_i intersects the boundary of V . Then, the entire path $\Sigma \cap F_i$ must be in the boundary of V or else there would be a fin F_j with $j < i$ intersecting F_i not at its endpoint, which would contradict Definition 1.2. In particular, ω must vanish along $\Sigma \cap F_i$.

Let $\tilde{\mathcal{A}}_i$ be the sheaf of piecewise linear functions on Δ whose divisors are supported on $\Sigma \cap F_i$. If we let $\tilde{\mathcal{D}}_i$ denote the quotient sheaf $\tilde{\mathcal{A}}_i/\mathbb{R}$, then we can give a global section $\tilde{\omega}_i$ of $\tilde{\mathcal{D}}_i$ defined piecewise such that $\tilde{\omega}_i$ vanishes on $\Delta \setminus F_i$ and it agrees with ω on $F_i \setminus \Sigma$. This defines a valid section of $\tilde{\mathcal{D}}_i$ because, as we noted, ω vanishes along $\Sigma \cap F_i$, and $\tilde{\omega}_i$ is clearly a section of \mathcal{D} away from $\Sigma \cap F_i$, on which we only require continuity. Consider the long exact sequence of cohomology associated with the quotient $\tilde{\mathcal{D}}_i$, analogous to (3):

$$0 \rightarrow H^0(\Delta, \mathbb{R}) \rightarrow H^0(\Delta, \tilde{\mathcal{A}}_i) \rightarrow H^0(\Delta, \tilde{\mathcal{D}}_i) \rightarrow H^1(\Delta, \mathbb{R}) \rightarrow$$

Since $\tilde{\omega}_i$ is only non-trivial on F_i , which is contractible, the image of $\tilde{\omega}_i$ in $H^1(\Delta, \mathbb{R})$ is trivial, so $\tilde{\omega}_i$ lifts to an element of $H^0(\Delta, \tilde{\mathcal{A}}_i)$, which we also denote by $\tilde{\omega}_i$ and we choose the representative such that $\tilde{\omega}_i$ is zero on Σ .

If $\tilde{\omega}_i$ is non-constant, then it must have a maximum value strictly greater than zero or minimum value strictly less than zero. Then, it would have its maximum or

minimum, respectively, on $F_i \setminus \Sigma$. We apply Proposition 2.8 to the linear function $\tilde{\omega}_i|_{F_i \setminus \Sigma}$ or to its negative to show that $\tilde{\omega}_i|_{F_i \setminus \Sigma}$ must be constant, and so $\tilde{\omega}_i$ is zero everywhere. Thus, ω is identically zero on F_i , and so we can expand V to include $F_i \setminus \Sigma$.

We've now shown that for each edge e of $\Sigma \cap F_i$, the section ω is zero on all but one simplex containing e , namely, the simplex in Σ on the other side from V . As in Definition 2.3, we let N_e^o denote the union interior of e with the interiors of the facets containing e . Since N_e^o is simply connected, we can lift $\omega|_{N_e^o}$ to a linear function ϕ on N_e^o . Then, by Lemma 2.4, ϕ is constant on N_e^o , and so ω vanishes on N_e^o . Therefore, we can further expand V to also include the interiors of all 2-simplices meeting F_i and the path $\Sigma \cap F_i$.

At the end, we will have that ω is zero on an open set V which contains the interior of every 2-simplex in Δ and since affine linear functions are continuous by definition, this means that ω is zero, which finishes the proof of the lemma. \square

We use Lemma 4.1 to prove the following strengthening of Theorem 1.4.

Theorem 4.2. *If Δ is a hyperbolic manifold with fins and ornaments, then there is no weak tropical surface, with Δ as its underlying topological space, and such that for every vertex v of Δ , M_v has at most one positive eigenvalue.*

Proof. Suppose that Δ is a weak tropical surface whose underlying Δ -complex is as in the theorem statement. We can assume that when decomposing Δ as in Definition 1.2, the subcomplex of ornaments \mathcal{O} is maximal, so that if we let Δ' denote the subcomplex consisting of just the manifold and fins, then Δ' is locally connected through codimension 1. Then, taking the restriction of the structure constants from Δ , we get that Δ' has the structure of a weak tropical surface, because the local structure around each edge of Δ' is unchanged. Moreover, at each vertex v of Δ' , the local intersection matrix M'_v is a block of the block diagonal matrix M_v for Δ . Therefore, M'_v also has at most one positive eigenvalue.

Next, we apply Lemma 3.2 to transform Δ' into a tropical surface Δ'' by gluing simplices onto edges of Δ' . Whenever we glue a simplex onto an edge e which is contained in one of the fins $F_i \subset \Delta'$, we can include that simplex in the fin, which remains contractible and its intersection with the manifold Σ is unchanged. If we glue a simplex onto an edge e contained in the manifold Σ , then the simplex forms a new fin F_{n+1} , numbered after all the other fins. Since $F_{n+1} \cap \Sigma$ is a single edge, the intersection of F_{n+1} with any other fin will be a subset of the endpoints of this edge. Thus, Δ'' is still a hyperbolic manifold with fins.

Finally, suppose that the manifold $\Sigma \subset \Delta''$ is not orientable. Then Σ has a 2-to-1 orientable cover, corresponding to an index 2 subgroup of its fundamental group [10, Proposition 3.25]. Since $\Sigma \subset \Delta''$ is a homotopy equivalence, the oriented cover extends to a cover of Δ'' , which we call $\widetilde{\Delta}''$. Each fin F_i of Δ'' is attached along a path of Σ , and so the preimage of F_i in $\widetilde{\Delta}''$ is two disjoint fins, which we number F_{2i-1} and F_{2i} to get a manifold with fins. Moreover, the Euler characteristic is multiplicative when taking covers, so $\chi(\widetilde{\Delta}'')$ is again negative. Therefore, we can replace Δ'' with $\widetilde{\Delta}''$ and so from now on we assume that the manifold $\Sigma \subset \Delta''$ is orientable.

By the classification of compact topological surfaces [7, Theorem 6.3], the Euler characteristic of a compact oriented surface is even, and so $\chi(\Delta'') \leq -2$. Therefore, $\dim_{\mathbb{R}} H^1(\Delta'', \mathbb{R}) = 2 - \chi(\Delta'') \geq 4$. By Theorem 2.6, the \mathbb{R} -span of the image of $H^0(\Delta'', \mathcal{D})$ has codimension at most 1 in $H^0(\Delta'', \mathbb{R})$, so the rank of $H^0(\Delta'', \mathcal{D})$ as an Abelian group is at least $\dim H^1(\Delta'', \mathbb{R}) - 1 \geq 3$. On the other hand, by Lemma 4.1, $H^0(\Delta'', \mathcal{D})$ is a subgroup of \mathbb{Z}^2 , and so a free Abelian group of rank at most 2. Therefore, we have a contradiction, so the weak tropical surface Δ cannot exist. \square

Proof of Theorem 4.3. Suppose \mathfrak{X} is a strict semistable degeneration whose dual complex Δ is a hyperbolic manifold with fins and ornaments. Then, by Proposition 3.1, Δ has the structure of a weak tropical complex such that the matrix M_v has at most one positive eigenvalue for every vertex v . However, by Theorem 4.2, such a weak tropical complex cannot exist, so we conclude that \mathfrak{X} cannot exist. \square

References

1. Baker, M.: Specialization of linear systems from curves to graphs. *Algebra Number Theory* **2**(6), 613–653 (2008). With an appendix by Brian Conrad
2. Berkovich, V.G.: Smooth p -adic analytic spaces are locally contractible. *Invent. Math.* **137**, 1–84 (1999)
3. Brown, M., Foster, T.: Rational connectivity and analytic contractibility. Preprint, [arXiv:1406.7312](https://arxiv.org/abs/1406.7312)
4. Cartwright, D.: Combinatorial tropical surfaces. Preprint, [arXiv:1506.02023](https://arxiv.org/abs/1506.02023)
5. Cartwright, D.: Tropical complexes. Preprint, [arXiv:1308.3813](https://arxiv.org/abs/1308.3813)
6. de Fernex, T., Kollár, J., Xu, C.: The dual complex of singularities. *Adv. Stud. Pure Math.* [in honor of Kawamata’s 60th birthday] (to appear). Preprint, [arXiv:1212.1675](https://arxiv.org/abs/1212.1675)
7. Gallier, J., Xu, D.: *A Guide to the Classification Theorem for Compact Surfaces*. Springer, Berlin (2013)
8. Gubler, W.: Local and canonical heights of subvarieties. *Ann. Sc. Norm. Super. Pisa Cl. Sci.* (5) **2**(4), 711–760 (2003)
9. Gubler, W., Rabinoff, J., Werner, A.: Skeletons and tropicalization. Preprint, [arXiv:1404.7044](https://arxiv.org/abs/1404.7044)
10. Hatcher, A.: *Algebraic Topology*. Cambridge University Press, Cambridge (2002)
11. Helm, D., Katz, E.: Monodromy filtrations and the topology of tropical varieties. *Can. J. Math.* **64**(4), 845–868 (2012)
12. Kapovich, M., Kollár, J.: Fundamental groups of links of isolated singularities. *J. Am. Math. Soc.* **27**(4), 929–952 (2014)
13. Kollár, J.: Simple normal crossing varieties with prescribed dual complex. *Algebra Geom.* **1**(1), 57–68 (2014)
14. Kulikov, V.S.: Degenerations of K3 surfaces and Enriques surfaces. *Izv. Akad. Nauk SSSR Ser. Mat.* **41**(5), 1008–1042 (1977)
15. Morrison, D.R.: Semistable degenerations of Enriques’ and hyperelliptic surfaces. *Duke Math. J.* **48**(1), 197–249 (1981)
16. Persson, U.: On degenerations of algebraic surfaces. *Mem. Am. Math. Soc.* 11(189), pp. xv+144 (1977)
17. Persson, U., Pinkham, H.: Degeneration of surfaces with trivial canonical bundle. *Ann. Math.* (2) **113**(1), 45–66 (1981)

Analytification and Tropicalization Over Non-archimedean Fields

Annette Werner

Abstract In this paper, we provide an overview of recent progress on the interplay between tropical geometry and non-archimedean analytic geometry in the sense of Berkovich. After briefly discussing results by Baker et al. (*Algebr. Geom.* **3**, 63–105 (2016); *Contemporary Mathematics*, vol 605, pp 93–121. American Mathematical Society, Providence, 2013) in the case of curves, we explain a result from Cueto et al. (*Math. Ann.* **360**, 391–437 (2014)) comparing the tropical Grassmannian of planes to the analytic Grassmannian. We also give an overview of most of the results in Gubler et al. (*Adv. Math.* **294**, 150–215 (2016)), where a general higher-dimensional theory is developed. In particular, we explain the construction of generalized skeleta in Gubler et al. (*Adv. Math.* **294**, 150–215 (2016)) which are polyhedral substructures of Berkovich spaces lending themselves to comparison with tropicalizations. We discuss the slope formula for the valuation of rational functions and explain two results on the comparison between polyhedral substructures of Berkovich spaces and tropicalizations.

Keywords Berkovich spaces • Tropical varieties • Skeletons • Slope formula • Faithful tropicalization

2010 MSC: 14G22, 14T05

1 Introduction

Tropical varieties are polyhedral images of varieties over non-archimedean fields. They are obtained by applying the valuation map to a set of toric coordinates. From the very beginning, analytic geometry was present in the systematic study

A. Werner (✉)

Institut für Mathematik, Goethe-Universität Frankfurt, Robert-Mayer-Strasse 8, 60325 Frankfurt, Germany

e-mail: werner@math.uni-frankfurt.de

of tropical varieties, e.g., in [16] where rigid analytic varieties are used. The theory of Berkovich spaces which leads to spaces with nicer topological properties than rigid varieties is even better suited to study tropicalizations.

A result of Payne [24] states that the Berkovich space associated with an algebraic variety is homeomorphic to the inverse limit of all tropicalizations in toric varieties. However, individual tropicalizations may fail to capture topological features of an analytic space. The present paper is an overview of recent results on the relationship between analytic spaces and tropicalizations. In particular, we address the question of whether a given tropicalization is contained in a Berkovich space as a combinatorial substructure.

A novel feature of Berkovich spaces compared to rigid analytic varieties is that they contain interesting piecewise linear combinatorial structures. In fact, Berkovich curves are, very roughly speaking, generalized graphs, where infinite ramifications along a dense set of points are allowed, see [4, Chap. 4], [1], and [3].

In higher dimensions, the structure of Berkovich analytic spaces is more involved, but still they often contain piecewise linear substructures as deformation retracts. This was a crucial tool in Berkovich's proof of local contractibility for smooth analytic spaces, see [5] and [6]. Berkovich constructs these piecewise linear substructures as the so-called skeleta of suitable models (or fibration of models). These skeleta basically capture the incidence structure of the irreducible components in the special fiber.

In dimension one, Baker, Payne, and Rabinoff studied the relationship between tropicalizations and subgraphs of Berkovich curves in [2] and [3]. Their results show that every finite subgraph of a Berkovich curve admits a faithful, i.e., homeomorphic and isometric tropicalization. They also prove that every tropicalization with tropical multiplicity one everywhere is isometric to a subgraph of the Berkovich curve.

As a first higher-dimensional example, the Grassmannian of planes was studied in [12]. The tropical Grassmannian of planes has an interesting combinatorial structure and is a moduli space for phylogenetic trees. It is shown in [12] that it is homeomorphic to a closed subset of the Berkovich analytic Grassmannian.

In [22], the higher-dimensional situation is analyzed from a general point of view. This approach is based on a generalized notion of a Berkovich skeleton which is associated with the datum of a semistable model plus a horizontal divisor. This naturally leads to unbounded skeleta and generalizes a well-known construction on curves, see [28] and [3].

For every rational function f with support in the fixed horizontal divisor it is shown in [22] that the "tropicalization" $\log |f|$ factors through a piecewise linear function on the generalized skeleton. This function satisfies a slope formula which is a kind of balancing condition around any one-codimensional polyhedral face. Moreover, for every generalized skeleton there exists a faithful tropicalization—where faithful in higher dimension refers to the preservation of the integral affine structures. In dimension one, this can be expressed via metrics. It is also proven that tropicalizations with tropical multiplicity one everywhere admit sections of the tropicalization map, which generalizes the above-mentioned result in [2] on curves.

The paper is organized as follows. In Sect. 2 we collect basic facts on Berkovich spaces and tropicalizations, giving references to the literature for proofs and more details. In Sect. 3 we briefly recall some of the results by Baker et al. [2] on curves. Section 4 starts with the definition and basic properties of the tropical Grassmannian. Theorem 4.1 claims the existence of a continuous section to the tropicalization map on the projective tropical Grassmannian. We give a sketch of the proof for the dense torus orbit, where some constructions are easier to explain than in the general case. In Sect. 5 we explain the construction of generalized skeleta from [22], and in Sect. 6 we investigate $\log |f|$ for rational functions f . In particular, Theorem 6.5 states the slope formula. Section 7 explains the faithful tropicalization results in higher dimension.

2 Berkovich Spaces and Tropicalizations

2.1 Notation and Conventions

A non-archimedean field is a field with a non-archimedean absolute value. Our ground field is a non-archimedean field which is complete with respect to its absolute value. Examples are the field \mathbb{Q}_p , which is the completion of \mathbb{Q} after the p -adic absolute value, finite extensions of \mathbb{Q}_p , and also the p -adic cousin \mathbb{C}_p of the complex numbers which is defined as the completion of the algebraic closure of \mathbb{Q}_p . The field of formal Laurent series $k((t))$ over an arbitrary base field k is another example. Besides, we can endow any field k with the trivial absolute value (which is one on all non-zero elements). A nice feature of Berkovich's general approach is that it also gives an interesting theory in the case of a trivially valued field. The reason behind this is that Berkovich geometry encompasses also points with values in transcendental field extensions—and those may well carry interesting non-trivial valuations.

Let K be a complete non-archimedean field. We write $K^\circ = \{x \in K : |x| \leq 1\}$ for the ring of integers in K , and $K^{\circ\circ} = \{x \in K : |x| < 1\}$ for the valuation ideal. The quotient $\tilde{K} = K^\circ / K^{\circ\circ}$ is the residue field of K . The valuation on K^\times associated with the absolute value is given by $v(x) = -\log |x|$. By $\Gamma = v(K^\times) \subset \mathbb{R}$ we denote the value group.

A variety over K is an irreducible, reduced, and separated scheme of finite type over K .

2.2 Berkovich Spaces

Let us briefly recall some basic results about Berkovich spaces. This theory was developed in the ground-breaking treatise [4]. The survey papers [11, 15], and [26] provide additional information. For background information on non-archimedean fields and Banach algebras and for an account of rigid analytic geometry see [8] and [7].

For $n \in \mathbb{N}$ and every n -tuple $r = (r_1, \dots, r_n)$ of positive real numbers we define the associated *Tate algebra* as

$$K\{r^{-1}x\} = \left\{ \sum_{I=(i_1, \dots, i_n) \in \mathbb{N}_0^n} a_I x^I : |a_I| r^I \rightarrow 0 \text{ as } |I| \rightarrow \infty \right\}.$$

Here we put $x = (x_1, \dots, x_n)$, and for $I = (i_1, \dots, i_n)$ we write $x^I = x^{i_1} \dots x^{i_n}$. Moreover, we define $|I| = i_1 + \dots + i_n$.

A Banach algebra A is called *K-affinoid*, if there exists a surjective K -algebra homomorphism $\alpha : K\{r^{-1}x\} \rightarrow A$ for some n and r such that the Banach norm on A is equivalent to the quotient seminorm induced by α . If such an epimorphism can be found with $r = (1, \dots, 1)$, then A is called *strictly K-affinoid*. In rigid analytic geometry only strictly K -affinoid algebras are considered (and called affinoid algebras).

The *Berkovich spectrum* $\mathcal{M}(A)$ of an affinoid algebra A is defined as the set of all bounded (by the Banach norm), multiplicative seminorms on A . It is endowed with the coarsest topology such that all evaluation maps on functions in A are continuous.

If x is a seminorm on an algebra A , and $f \in A$, we follow the usual notational convention and write $|f(x)|$ for $x(f)$, i.e., for the real number which we get by evaluating the seminorm x on f .

The *Shilov boundary* of a Berkovich spectrum $\mathcal{M}(A)$ of a K -affinoid algebra A is the unique minimal subset Γ (with respect to inclusion) such that for every $f \in A$ the evaluation map $\mathcal{M}(A) \rightarrow \mathbb{R}_{\geq 0}$ given by $x \mapsto |f(x)|$ attains its maximum on Γ . Hence for every $f \in A$ there exists a point z in the Shilov boundary such that $|f(z)| \geq |f(x)|$ for all $x \in \mathcal{M}(A)$. The Shilov boundary of $\mathcal{M}(A)$ exists and is a finite set by Berkovich [4, Corollary 2.4.5].

The Berkovich spectrum of a K -affinoid algebra carries a sheaf of analytic functions which we will not define here. Some care is needed since, very roughly speaking, not all coverings are suitable for glueing analytic functions. Analytic spaces over K are ringed spaces which are locally modelled on Berkovich spectra of affinoid algebras.

There is a *GAGA-functor*, associating to every K -scheme X locally of finite type an analytic space X^{an} over K . It has the property that X is connected, separated over K , or proper over K , respectively, if and only if the topological space X^{an} is arcwise connected, Hausdorff, or compact, respectively. If $X = \text{Spec}R$ is affine, the topological space X^{an} can be identified with the set of all multiplicative seminorms on the coordinate ring R extending the absolute value on K . This space is endowed with the coarsest topology such that for all $f \in R$ the evaluation map on f is continuous.

With the help of the GAGA-functor we associate analytic spaces to algebraic varieties. Another way of obtaining analytic spaces is via *admissible formal schemes*. An admissible formal scheme over K° is, roughly speaking, a formal scheme over K° such that the formal affine building blocks are given by K° -flat

algebras of the form $K^\circ\{x_1, \dots, x_n\}/\mathfrak{a}$ for a finitely generated ideal \mathfrak{a} in the ring $K^\circ\{x_1, \dots, x_n\} = \{\sum_{I \in \mathbb{N}_0^n} a_I x^I : a_I \in K^\circ \text{ and } |a_I| \rightarrow 0 \text{ as } |I| \rightarrow \infty\}$. For a precise definition see [7] or [11].

An admissible formal scheme \mathcal{X} has a natural *analytic generic fiber* \mathcal{X}_η . On the formal affine building block given by $K^\circ\{x_1, \dots, x_n\}/\mathfrak{a}$ the analytic generic fiber is the Berkovich spectrum of $K\{x_1, \dots, x_n\}/\mathfrak{a}K\{x_1, \dots, x_n\}$.

If we start with a K° -scheme \mathcal{X} , which is separated, flat, and of finite presentation, its completion with respect to any element $\pi \in K^{\circ\circ} \setminus \{0\}$ is an admissible formal scheme \mathcal{X} . This has an analytic generic fiber \mathcal{X}_η . On the other hand, we can apply the GAGA-functor to the algebraic generic fiber $X = \mathcal{X} \otimes_{K^\circ} K$ and get another analytic space X^{an} . They are connected by a morphism $\mathcal{X}_\eta \hookrightarrow X^{\text{an}}$, which is an isomorphism if \mathcal{X} is proper over K° . In the basic example $\mathcal{X} = \text{Spec}K^\circ[x]$ we get the natural inclusion $\mathcal{M}(K\{x\}) \hookrightarrow (\mathbb{A}_K^1)^{\text{an}}$, whose image is the set of all multiplicative seminorms on the polynomial ring $K[x]$ which extend the absolute value on K and whose value on x is bounded by 1.

2.3 Tropicalization

We start by considering a split torus $T = \text{Spec}K[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$. Evaluating seminorms on characters we get a map $\text{trop} : T^{\text{an}} \rightarrow \mathbb{R}^n$ on the associated Berkovich space, which is given by

$$\text{trop}(p) = (-\log |x_1(p)|, \dots, -\log |x_n(p)|).$$

Here we are using the valuation map to define tropicalizations. In some papers, tropicalizations are defined via the negative valuation map, i.e., the logarithmic absolute value of the coordinates.

If X is a variety over K together with a closed embedding $\varphi : X \hookrightarrow T$, we consider the composition

$$\text{trop}_\varphi : X^{\text{an}} \xrightarrow{\varphi^{\text{an}}} T^{\text{an}} \xrightarrow{\text{trop}} \mathbb{R}^n \tag{1}$$

of the tropicalization map with the embedding φ . Note that the map trop_φ is continuous. Its image $\text{Trop}_\varphi(X) = \text{trop}_\varphi(X^{\text{an}})$ is the support of a polyhedral complex Σ in \mathbb{R}^n . This complex is integral Γ -affine in the sense of Sect. 5.1. It has pure dimension $d = \dim(X)$. For a nice introductory text on tropical geometry see the textbook [23]. The survey paper [21] is also very useful.

Note that for all $\omega \in \text{Trop}_\varphi(X)$ the preimage $\text{trop}_\varphi^{-1}(\{\omega\})$ can be identified with the Berkovich spectrum of an affinoid algebra, see [20, Proposition 4.1].

For every point $\omega \in \text{Trop}_\varphi(X)$ there is an associated *initial degeneration*, which is the special fiber of a model of X over the valuation ring in a suitable non-archimedean extension field of K . The geometric number of irreducible components

of this special fiber is the *tropical multiplicity* of ω . There is a balancing formula for tropicalizations involving tropical multiplicities on all maximal-dimensional faces of $\text{Trop}_\varphi(X)$ around a fixed codimension one face, see [23, Chap. 3.4].

Let \overline{K} be the algebraic closure of K . It can be endowed with an absolute value extending the one on K . By means of the coordinates x_1, \dots, x_n we can identify $T(\overline{K}) = \overline{K}^n$. Then we can define a natural tropicalization map $\text{trop} : T(\overline{K}) \rightarrow \mathbb{R}^n$ by $\text{trop}(t_1, \dots, t_n) = (-\log |t_1|, \dots, -\log |t_n|)$. Note that every point $t = (t_1, \dots, t_n)$ in $T(\overline{K})$ gives rise to the multiplicative seminorm $f \mapsto |f(t)|$ on the coordinate ring $K[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$, which is an element in T^{an} . Hence the tropicalization map on $T(\overline{K})$ is induced by the map trop on T^{an} .

We could define the tropicalization $\text{Trop}_\varphi(X)$ without recourse to Berkovich spaces. If the absolute value on K is non-trivial, the tropicalization $\text{Trop}_\varphi(X)$ is equal to the closure of the image of the map

$$X(\overline{K}) \xrightarrow{\varphi} T(\overline{K}) \xrightarrow{\text{trop}} \mathbb{R}^n.$$

Considering a tropical variety as the image of an analytic space makes some topological considerations easier. For example, in this way it is evident that the tropicalization of a connected variety is connected since it is a continuous image of a connected space.

Let Y be a toric variety associated with the fan Δ in $N_{\mathbb{R}}$, where N is the cocharacter group of the dense torus T . Then Δ defines a natural partial compactification $N_{\mathbb{R}}^\Delta$ (due to Kajiwara and Payne) of the space $N_{\mathbb{R}}$. Roughly speaking, we compactify cones by dual spaces. For details see [24], Sect. 3. There is also a natural tropicalization map

$$\text{trop} : Y^{\text{an}} \rightarrow N_{\mathbb{R}}^\Delta,$$

which extends the tropicalization map on the dense torus. As an example, we consider $Y = \mathbb{P}_K^n$ with its dense torus $\mathbb{G}_{m,K}^n$ and $N = \mathbb{Z}^n$. Put $\overline{\mathbb{R}} = \mathbb{R} \cup \{\infty\}$ and endow it with the natural topology such that half-open intervals $]a, \infty]$ form a basis of open neighborhoods of ∞ . Then the associated compactification of $N_{\mathbb{R}} = \mathbb{R}^n$ is the tropical projective space

$$\mathbb{TP}^n = (\overline{\mathbb{R}}^{n+1} \setminus \{(\infty, \dots, \infty)\}) / \mathbb{R}(1, \dots, 1)$$

which is endowed with the product-quotient topology. The Berkovich projective space $(\mathbb{P}_K^n)^{\text{an}}$ can be described as the set of equivalence classes of elements in $(\mathbb{A}_K^{n+1})^{\text{an}} \setminus \{0\}$ with respect to the following equivalence relation: Let x and y be points in $(\mathbb{A}_K^{n+1})^{\text{an}} \setminus \{0\}$, i.e., multiplicative seminorms on $K[x_0, \dots, x_n]$ extending the absolute value on K . They are equivalent if and only if there exists a constant $c > 0$ such that for every homogeneous polynomial f of degree d we have $|f(y)| = c^d |f(x)|$.

We can describe the tropicalization map on the toric variety \mathbb{P}_K^n as follows:

$$\text{trop} : (\mathbb{P}_K^n)^{\text{an}} \rightarrow \mathbb{T}\mathbb{P}^n \tag{2}$$

maps the class of the seminorm p on $K[x_0, \dots, x_n]$ to the class $(-\log |x_0(p)|, \dots, -\log |x_n(p)|) + \mathbb{R}(1, \dots, 1)$ in tropical projective space.

An important result about the relation between tropical and analytic geometry due to Payne says that for every quasi-projective variety X over K the associated analytic space X^{an} is homeomorphic to the inverse limit over all $\text{Trop}_\varphi(X)$, where the limit runs over all closed embeddings $\varphi : X \hookrightarrow Y$ in a quasi-projective toric variety Y (see [24, Theorem 4.2]). For a generalization which omits the quasi-projectivity hypothesis see [17].

Hence tropicalizations are combinatorial images of analytic spaces, which recover the full analytic space in the projective limit. Individual tropicalizations however may not faithfully depict all topological features of the analytic space. We will discuss a basic example at the beginning of the next section.

3 The Case of Curves

Let K be a field which is algebraically closed and complete with respect to a non-archimedean, non-trivial absolute value, and let X be a smooth curve over K .

As an example, let us consider an elliptic curve E in Weierstrass form $y^2 = x^3 + ax + b$ over K . The affine curve E has a natural embedding in \mathbb{A}_K^2 . The intersection E_0 of E with $\mathbb{G}_{m,K}^2$ has a natural tropicalization in \mathbb{R}^2 , which is one of the trees depicted in the next figure.

If however E is a Tate curve, then the analytic space E^{an} contains a circle, i.e., it has topological genus 1. Hence the Weierstrass tropicalization does not faithfully depict the topology of the analytic space.

As stated before, if X is a smooth curve over K , the Berkovich space X^{an} is some kind of generalized graph, allowing infinite ramification along a dense set of points. In particular, it makes sense to talk about its leaves. The space of non-leaves $H_0(X^{\text{an}})$ in the Berkovich analytification X^{an} admits a natural metric which is defined via semistable models, see [3, Sect. 5.3]. Any tropicalization of X with respect to a closed embedding in a torus carries a natural metric which is locally given by lattice length on each edge (and globally by the shortest paths). Another problem in the case of the Weierstraß tropicalizations from Fig. 1 is the fact that tropical multiplicities $\neq 1$ are present. This may indicate that the tropicalization map will not be isometric on a subgraph, see [2, Corollary 5.9].

It is explained in [2, Theorem 6.2], how one can construct better tropicalizations of the Tate curve.

Based on a detailed study of the structure of analytic curves there are the following two important comparison theorems in [2].




Fig. 1 Tropicalization of an elliptic curve in Weierstrass form. The *left-hand side* shows the case $3v(a) \geq 2v(b)$ and the *right-hand side* the case $3v(a) < 2v(b)$. The numbers on the edges indicate the tropical multiplicities. If no number is given, the multiplicity is one

Theorem 3.1 ([2, Theorem 5.20]). *Let Γ be any finite subgraph of X^{an} for a smooth, complete curve X over K . Then there exists a closed immersion $\varphi : X \hookrightarrow Y$ into a toric variety Y with dense torus T such that the restriction $\varphi_0 : X_0 = X \cap \varphi^{-1}(T) \hookrightarrow T$ induces a tropicalization map $\text{trop}_{\varphi_0} : X_0^{\text{an}} \rightarrow \mathbb{R}^n$ which maps Γ homeomorphically and isometrically onto its image.*

Theorem 3.2 ([2, Theorem 5.24]). *Consider a smooth curve X_0 over K and a closed immersion $\varphi : X_0 \hookrightarrow T$ into a split torus with associated tropical variety $\text{Trop}_{\varphi_0}(X_0)$. If Γ' is a compact connected subset of $\text{Trop}_{\varphi_0}(X_0)$ which has tropical multiplicity one everywhere, then there exists a unique closed subset Γ in $H_0(X_0^{\text{an}})$ mapping homeomorphically onto Γ' , and this homeomorphism is in fact an isometry.*

We will discuss higher-dimensional generalizations of these theorems in Sect. 7.

4 Tropical Grassmannians

4.1 The Setting

In view of the results [2] for curves, it is a natural question whether they can be generalized to varieties of higher dimensions. As a first example, the tropical Grassmannian of planes was studied in [12]. In this section we discuss the main theorem of this paper.

For natural numbers $d \leq n$ we denote by $\text{Gr}(d, n)$ the Grassmannian of d -dimensional subspaces of n -space. The tropical Grassmannian is defined as the tropicalization of $\text{Gr}(d, n)$ with respect to the Plücker embedding

$\varphi : \text{Gr}(d, n) \rightarrow \mathbb{P}^{\binom{n}{d}-1}$. Recall that the Plücker embedding maps the point corresponding to the d -dimensional subspace W to the line in $\binom{n}{d}$ -space given by the d th exterior power of W .

During this section we will deviate from the exposition in Sect. 2.3 and consider tropicalizations with respect to the negative valuation map. The choice between $(\min, +)$ and $(\max, +)$ tropical geometry is always a difficult one.

Hence our tropical projective space is

$$\mathbb{TP}^{\binom{n}{d}-1} = ((\mathbb{R} \cup \{-\infty\})^{\binom{n}{d}} \setminus \{(-\infty, \dots, -\infty)\}) / \mathbb{R}(1, \dots, 1),$$

and the tropical Grassmannian $\mathcal{T}\text{Gr}(d, n) = \text{Trop}_\varphi(\text{Gr}(d, n))$ is defined as the image of $\text{trop} \circ \varphi^{\text{an}} : \text{Gr}(d, n)^{\text{an}} \rightarrow \mathbb{TP}^{\binom{n}{d}-1}$, where trop is given by the map

$$p \mapsto (\log |x_0(p)|, \dots, \log |x_{\binom{n}{d}}(p)|) + \mathbb{R}(1, \dots, 1)$$

on the analytic projective space, see (2).

Tropical Grassmannians were first studied by Speyer and Sturmfels [25] who focused on the toric part $\text{Gr}_0(d, n) = \text{Gr}(d, n) \cap \varphi^{-1}(\mathbb{G}_{m,K}^{\binom{n}{d}-1})$ which is embedded in the torus $\mathbb{G}_{m,K}^{\binom{n}{d}-1}$ via φ . The tropicalization $\mathcal{T}\text{Gr}_0(d, n) = \text{Trop}_\varphi(\text{Gr}_0(d, n))$ (which is called $\mathcal{G}'_{d,n}$ in [25]) is a fan of dimension $d(n - d)$ containing a linear space of dimension $n - 1$, see [25, Sect. 3].

Moreover, if $d = 2$, Speyer and Sturmfels proved that $\mathcal{T}\text{Gr}_0(2, n)$ can be identified with the space of phylogenetic trees, see [25, Sect. 4]. For our purposes, a phylogenetic tree on n leaves is a pair (T, ω) , where T is a finite combinatorial tree with no degree-two vertices together with a labelling of its leaves in bijection with $\{1, \dots, n\}$, and ω is a real-valued function on the set of edges of T . The tree T is called the combinatorial type of the phylogenetic tree (T, ω) . For every phylogenetic tree (T, ω) and all $i \neq j$ in $\{1, \dots, n\}$ we denote by x_{ij} the sum of the weights along the uniquely determined path from leaf i to leaf j . This tree-distance function satisfies the four-point-condition, which states that for all pairwise distinct indices i, j, k, l in $\{1, \dots, n\}$ the maximum among

$$x_{ij} + x_{kl}, \quad x_{ik} + x_{jl}, \quad x_{il} + x_{jk}$$

is attained at least twice.

4.2 A Section of the Tropicalization Map

For the rest of this section we will always consider the case $d = 2$. In [12], we investigate the full projective Grassmannian $\mathcal{T}\text{Gr}(2, n) = \text{Trop}_\varphi \text{Gr}(2, n)$ in tropical projective space $\mathbb{TP}^{\binom{n}{2}-1}$. The main result of this paper is the following theorem:

Theorem 4.1 ([12, Theorem 1.1]). *There exists a continuous section $\sigma: \mathcal{T}\text{Gr}(2, n) \rightarrow \text{Gr}(2, n)^{\text{an}}$ of the tropicalization map $\text{trop} \circ \varphi^{\text{an}}: \text{Gr}(2, n)^{\text{an}} \rightarrow \mathcal{T}\text{Gr}(2, n)$. Hence, the tropical Grassmannian $\mathcal{T}\text{Gr}(2, n)$ is homeomorphic to a closed subset of the Berkovich analytic space $\text{Gr}(2, n)^{\text{an}}$.*

Note that we are not considering semistable models or their Berkovich skeleta in this approach.

The first idea one might have is to look at the big open cells of the Grassmannians. Since $d = 2$, the coordinates of the ambient projective space are indexed by the two element subsets $\{i, j\}$ of $\{1, \dots, n\}$. If $i < j$, we write p_{ij} for this coordinate, and we put $p_{ji} = -p_{ij}$. We intersect the Grassmannian with the standard open affine covering of projective space and get open affine subvarieties

$$U_{ij} = \text{Gr}(2, n) \cap \varphi^{-1}\{p_{ij} \neq 0\}.$$

The affine Plücker coordinates are $u_{kl} = p_{kl}/p_{ij}$. Since the Plücker ideal is generated by the relations $p_{ij}p_{kl} - p_{ik}p_{jl} + p_{il}p_{jk} = 0$ for $\{i, j, k, l\}$ running over the four-element subsets of $\{1, \dots, n\}$, the subvariety U_{ij} can be identified with $\mathbb{A}_K^{n(n-2)}$ by means of the coordinates $u_{ik} = p_{ik}/p_{ij}$ and $u_{jk} = p_{jk}/p_{ij}$ for all k different from i and j .

Recall that $(\mathbb{A}_K^{n(n-2)})^{\text{an}}$ is the set of all multiplicative seminorms on $K[u_{ik}, u_{jk} : k \notin \{i, j\}]$ which extend the absolute value on K . As explained in Sect. 2, the toric variety $\mathbb{A}_K^{n(n-2)}$ has a natural tropicalization. With the sign conventions in the present section, it is given by

$$\begin{aligned} \text{trop} : (\mathbb{A}_K^{n(n-2)})^{\text{an}} &\rightarrow (\mathbb{R} \cup \{-\infty\})^{2(n-2)} \\ p &\mapsto (\log |u_{ik}(p)|, \log |u_{jk}(p)|)_{k \notin \{i, j\}} \end{aligned}$$

This induces the tropicalization map

$$\text{trop} : U_{ij}^{\text{an}} \rightarrow (\mathbb{R} \cup \{-\infty\})^{2(n-2)}.$$

Now the tropicalization map on an affine space $(\mathbb{A}_K^N)^{\text{an}}$ has a natural section, which is given by the map $\delta : (\mathbb{R} \cup \{-\infty\})^N \rightarrow (\mathbb{A}_K^N)^{\text{an}}$, mapping a point $r = (r_1, \dots, r_N) \in (\mathbb{R} \cup \{-\infty\})^N$ to the seminorm

$$\delta(r): K[x_1, \dots, x_N] \longrightarrow \mathbb{R}_{\geq 0} \quad \left| \sum_{\alpha \in \mathbb{N}_0^N} c_\alpha x^\alpha (\delta(r)) \right| = \max_\alpha \left\{ |c_\alpha| \prod_{i=1}^N \exp(r_i \alpha_i) \right\}. \tag{3}$$

Here we put $\exp((-\infty)\alpha) = 0$ for all $\alpha \in \mathbb{N}$. Note that the restriction of δ to \mathbb{R}^N is a map from \mathbb{R}^N to the torus $(\mathbb{G}_m^N)^{\text{an}}$. We call the image of δ on $(\mathbb{R} \cup \{-\infty\})^N$ the standard skeleton of $(\mathbb{A}_K^N)^{\text{an}}$, and the image of $\delta|_{\mathbb{R}^N}$ the standard skeleton of the torus $(\mathbb{G}_m^N)^{\text{an}}$.

However, the map δ for $N = n(n - 2)$ does not in general provide a section of the tropicalization map on the whole of U_{ij} . Consider $n \geq 4$, and let $x = (x_{kl})_{kl}$ be a point in the tropical Grassmannian $\mathcal{TGr}(2, n)$ which lies in the image of U_{ij} under the tropicalization map. Let ω be the projection to the affine coordinates $\omega = (x_{ik} - x_{ij}, x_{jk} - x_{ij})_{k \notin \{i, j\}} \in \mathbb{R}^{2(n-2)}$. Then $\delta(\omega)$ is a point in the Berkovich space U_{ij}^{an} with $\log |u_{ik}(\delta(\omega))| = x_{ik} - x_{ij}$ and $\log |u_{jk}(\delta(\omega))| = x_{jk} - x_{ij}$. Since $u_{kl} := p_{kl}/p_{ij} = u_{ik}u_{jl} - u_{il}u_{jk}$ by the Plücker relations, the definition of δ gives us $\log |u_{kl}(\delta(\omega))| = \max\{x_{ik} + x_{jl}, x_{il} + x_{jk}\} - 2x_{ij}$. Hence δ can only provide a section of the tropicalization map if this maximum is equal to $x_{kl} + x_{ij}$. This may fail if the labelled tree T has the wrong shape, e.g., if it looks like this:



The strategy for the proof of Theorem 4.1 is to compare the tropical Grassmannian with a standard skeleton of an affine space on a smaller piece of the tropicalization, namely, on the part consisting of phylogenetic trees such that the underlying combinatorial tree has the right shape. If we want to make this precise, the definition of these smaller pieces is quite involved. A considerable part of the difficulties is due to the fact that we take into account the boundary strata of the Grassmannian in projective space. In order to explain the general strategy we will from now on restrict our attention to the torus part of the tropical Grassmannian. The general case can be found in [12].

4.3 Sketch of Proof in the Dense Torus Orbit

In this section we explain the proof of Theorem 4.1 for the subset $\mathcal{TGr}_0(2, n)$ of the tropical Grassmannian $\mathcal{TGr}(2, n)$. Recall that $\mathcal{TGr}_0(2, n)$ is the tropicalization of the dense open subset $\text{Gr}_0(2, n)$ of the Grassmannian which is mapped to the torus via the Plücker map.

We fix a pair ij as above and work in the big open cell U_{ij} . The coordinate ring R_{ij} of U_{ij} is a polynomial ring $K[u_{ik}, u_{jk} : k \notin \{i, j\}]$ in $2(n - 2)$ variables. The other Plücker coordinates are expressed as $u_{kl} = u_{ik}u_{jl} - u_{il}u_{jk}$ in R_{ij} . Note that the affine variety $\text{Gr}_0(2, n)$ is contained in U_{ij} . The coordinate ring of $\text{Gr}_0(2, n)$ is equal to the localization of R_{ij} after the multiplicative subset generated by all u_{kl} for $\{k, l\} \neq \{i, j\}$.

We also fix a labelled tree T with n leaves $1, \dots, n$, and arrange T as in Fig. 2 with subtrees T_1, \dots, T_r .

Fig. 2 Arrangement of T





Fig. 3 Cherry property



Definition 4.2. Let \preceq be a partial order on the set $\{1, \dots, n\} \setminus \{i, j\}$. We write $k < l$ if $k \preceq l$ and $k \neq l$. Then \preceq has the *cherry property* on T (see Fig. 3) with respect to i and j if the following conditions hold:

- (i) Two leaves of different subtrees T_a and T_b for $a, b \in \{1, \dots, r\}$ as in Fig. 2 cannot be compared by \preceq .
- (ii) The partial order \preceq restricts to a total order on the leaf set of each T_a , $a = 1, \dots, r$.
- (iii) If $k < l < m$, then either $\{k, l\}$ or $\{l, m\}$ is a *cherry of the quartet* $\{i, k, l, m\}$, i.e., the subtree given by this quartet of leaves contains a node which is adjacent to both elements of the cherry:

An induction argument shows the following lemma:

Lemma 4.3 ([12, Lemma 4.7]). Fix a pair of indices i, j , and let T be a tree on n labelled leaves. Then, there exists a partial order \preceq on the set $\{i, \dots, n\} \setminus \{i, j\}$ that has the cherry property on T with respect to i and j .

For an example of the inductive construction of such a partial order see [12], Fig. 3.

The leaves in each subtree T_a are totally ordered, say as $s_1 < s_2 < \dots < s_p$, if T_a contains $p = p(a)$ leaves. We consider the variable set $I_a = \{u_{is_1}, \dots, u_{is_p}\} \cup \{u_{js_1}, u_{s_1s_2}, \dots, u_{s_{p-1}s_p}\}$. Then $I = I_1 \cup \dots \cup I_r$ is a set of $2(n - 2)$ affine Plücker coordinates of the form $u_{kl} \in R_{ij}$.

We can successively reconstruct the variables of the form u_{jl} which are not contained in I as follows: In the tree T_a with leaves $s_1 < s_2 < \dots < s_p$, we have $u_{s_1s_2} = u_{is_1}u_{js_2} - u_{is_2}u_{js_1}$, hence

$$u_{js_2} = u_{is_1}^{-1}(u_{s_1s_2} + u_{is_2}u_{js_1}).$$

The right-hand side is an expression in the variables contained in the coordinate set I with u_{is_1} inverted. Now we use the relation $u_{s_2s_3} = u_{is_2}u_{js_3} - u_{is_3}u_{js_2}$ to express $u_{js_3} = u_{is_2}^{-1}(u_{s_2s_3} + u_{is_3}u_{js_2})$. Plugging in the expression for u_{js_2} we can write u_{js_3} as a polynomial in the variables in I plus all u_{ik}^{-1} .

Proceeding by induction, we find for all $m \neq i, j$ that

$$u_{jm} \in K[u_{kl} : u_{kl} \in I][u_{ik}^{-1} : k \neq i, j]$$

and hence

$$K[u_{kl} : u_{kl} \in I] \subset R_{ij} \subset K[u_{kl} : u_{kl} \in I][u_{ik}^{-1} : k \neq i, j]. \tag{4}$$

This shows that the variable set I generates the function field $\text{Quot}(R_{ij})$ of U_{ij} . Recall that the coordinate ring of $\text{Gr}_0(2, n)$ is $K[\text{Gr}_0(2, n)] = S^{-1}R_{ij}$, where S is the multiplicative subset of R_{ij} generated by all u_{kl} for $\{kl\} \neq \{ij\}$. By the previous result, $K[\text{Gr}_0(2, n)]$ is equal to the localization of the Laurent polynomial ring $K[u_{kl}^{\pm} : u_{kl} \in I]$ after the multiplicative subset generated by *all* u_{kl} expressed as Laurent polynomials in the coordinates contained in I .

Definition 4.4. Let \mathcal{C}_T be the cone in $\mathcal{TGr}_0(2, n)$ whose interior corresponds to the phylogenetic trees with underlying tree T . Let $x = (x_{kl})_{kl} + \mathbb{R}(1, \dots, 1)$ be a point in $\mathcal{C}_T \subset \mathcal{TGr}_0(2, n)$. We associate to it a point $\sigma_T^{ij}(x)$ in $\text{Gr}_0(2, n)^{\text{an}}$, i.e., a multiplicative seminorm on the coordinate ring of $\text{Gr}_0(2, n)$, as follows. For every Laurent polynomial $f = \sum_{\alpha} c_{\alpha} u^{\alpha} \in K[u_{kl}^{\pm} : u_{kl} \in I]$ (where α runs over \mathbb{Z}^I) we put

$$|f(\sigma_T^{ij}(x))| = \max_{\alpha} \{ |c_{\alpha}| \prod_{u_{kl} \in I} \exp(\alpha_{kl}(x_{kl} - x_{ij})) \}, \quad \text{where } \alpha = (\alpha_{kl})_{u_{kl} \in I}.$$

This defines a multiplicative norm on $K[u_{kl}^{\pm} : \{kl\} \in I]$ which has a unique extension to a multiplicative norm on the localization $K[\text{Gr}_0(2, n)]$. Let $\sigma_T^{ij}(x)$ be the resulting point in the Berkovich space $\text{Gr}_0(2, n)^{\text{an}}$.

Since for every $f \in K[\text{Gr}_0(2, n)]$ the evaluation map on \mathcal{C}_T given by

$$x \mapsto |f(\sigma_T^{ij}(x))|$$

is continuous, we have constructed a continuous map

$$\sigma_T^{ij} : \mathcal{C}_T \rightarrow \text{Gr}_0(2, n)^{\text{an}}.$$

We want to show that it is a section of the tropicalization map $\text{trop} \circ \varphi^{\text{an}} : \text{Gr}_0(2, n)^{\text{an}} \rightarrow \mathcal{TGr}_0(2, n)$ on $\text{trop}^{-1}(\mathcal{C}_T)$, which amounts to checking that

$$\log |u_{kl}(\sigma_T^{ij}(x))| = x_{kl} - x_{ij} \text{ for all } \{kl\} \neq \{ij\} \text{ and for all } x \in \mathcal{C}_T.$$

For variables u_{kl} in I this is clear from the definition of σ_T^{ij} . Hence it holds in particular for all indices of the form $\{ik\}$ for $k \notin \{i, j\}$. In order to check this fact for the other indices, recall the definition of I after Lemma 4.3. Since $u_{js_2} = u_{is_1}^{-1}(u_{s_1s_2} + u_{is_2}u_{js_1})$, we find

$$\log |u_{js_2}(\sigma_T^{ij}(x))| = \max\{-x_{is_1} + x_{s_1s_2}, -x_{is_1} + x_{is_2} + x_{js_1} - x_{ij}\}.$$

Since s_1 and s_2 are in the same subtree (see Fig. 2), we find that $x_{is_1} + x_{js_2} = x_{is_2} + x_{js_1} \geq x_{ij} + x_{s_1s_2}$, which implies $\log |u_{js_2}(\sigma_T^{ij}(x))| = x_{js_2} - x_{ij}$. Inductively, we can show in this way our claim for all indices of the form $\{jk\}$. If we consider an index of the form $\{kl\}$ where $k, l \notin \{i, j\}$, we have $u_{kl} = u_{ik}u_{jl} - u_{il}u_{jk}$. If k is a leaf in the subtree T_a and l is a leaf in the subtree T_b for $a < b$, we find $x_{kl} + x_{ij} = x_{il} + x_{jk} > x_{ik} + x_{jl}$. Hence the non-archimedean triangle inequality gives $\log |u_{kl}(\sigma_T^{ij}(x))| = \log |(u_{il}u_{jk})(\sigma_T^{ij}(x))| = x_{kl} - x_{ij}$. If k and l are leaves in the same subtree T_a , we may assume that $k < l$. If k is the predecessor of l in the total ordering \leq restricted to T_a , we are done, since then u_{kl} is contained in I . If not, we let m be the predecessor of l , so that $k < m < l$. The Plücker relations give $u_{im}u_{kl} = u_{ik}u_{ml} + u_{il}u_{km}$, hence

$$u_{kl} = u_{im}^{-1}(u_{ik}u_{ml} + u_{il}u_{km}).$$

Note that all variables on the right-hand side except possibly u_{km} are contained in I . Hence we can calculate $\log |u_{kl}(\sigma_T^{ij}(x))|$ by developing u_{km} in a Laurent series in the variables in I . The resulting Laurent series does not allow any cancellation between the variables. Now we can apply the cherry property for the ordering \leq , which says that $\{k, m\}$ or $\{m, l\}$ is a cherry for the quartet $\{i, k, m, l\}$. Hence we have

$$x_{il} + x_{km} \leq x_{im} + x_{kl} = x_{ik} + x_{ml}$$

or

$$x_{ik} + x_{ml} \leq x_{im} + x_{kl} = x_{il} + x_{km}.$$

In both cases, one can check directly that our claim holds.

In fact, the section σ_T^{ij} has a more conceptual description. Let $\psi : \text{Gr}_0(2, n) \rightarrow \text{Spec}K[u_{kl}^{\pm} : u_{kl} \in I] = \mathbb{G}_m^{2(n-2)}$ be the open embedding induced by (4), and let Σ be the standard skeleton of $(\mathbb{G}_m^{2(n-2)})^{\text{an}}$ as in (3). It is contained in the open analytic subvariety $\text{Gr}_0(2, n)^{\text{an}}$, since it consists entirely of norms and does therefore not meet a closed subvariety of strictly lower dimension. Then the map σ_T^{ij} is the composition of the projection from \mathcal{C}_T to the coordinates $(x_{kl} - x_{ij})_{u_{kl} \in I} \in \mathbb{R}^{2(n-2)} = \Sigma$ followed by the inclusion of Σ in $\text{Gr}_0(2, n)^{\text{an}}$.

Recall from Sect. 2.3 that for every $x \in \mathcal{T}\text{Gr}_0(2, n)$ the preimage $\text{trop}_{\varphi}^{-1}(x)$ under the tropicalization map is the Berkovich spectrum $\mathcal{M}(A_x)$ of an affinoid algebra A_x . It follows from our construction that $\sigma_T^{ij}(x)$ is the unique Shilov boundary point of A_x . We can formulate this fact explicitly as follows.

Lemma 4.5 ([12, Lemma 4.17]). *For every $x \in \mathcal{C}_T$ the seminorm $\sigma_T^{ij}(x)$ constructed above has the following maximality property: For every f in the coordinate ring $K[\text{Gr}_0(2, n)]$ and every seminorm $p \in \text{trop}_{\varphi}^{-1}(x) \subset \text{Gr}_0(2, n)^{\text{an}}$ we have*

$$|f(p)| \leq |f(\sigma_T^{ij}(x))|.$$

As an immediate consequence we see that $\sigma_T^{ij}(x)$ does not depend on the choice of the partial ordering \preceq or on the pair ij , and that $\sigma_T^{ij}(x) = \sigma_{T'}^{ij}(x)$ if x is contained in the intersection $\mathcal{C}_T \cap \mathcal{C}_{T'}$. Hence we can patch these maps together and get a well-defined section $\sigma : \mathcal{T}\text{Gr}_0(2, n) \rightarrow \text{Gr}_0(2, n)^{\text{an}}$ of the tropicalization map. It follows from the construction that it is continuous.

This proves Theorem 4.1 on the dense torus orbit $\text{Gr}_0(2, n)$. In order to define a section of the tropicalization map also on the boundary strata of $\text{Gr}(2, n)$, we follow the same strategy of constructing an index set I such that the associated Plücker variables generate the function field of the Grassmannian. This construction is more involved, see [12, Sect. 4] for details. The most important problem here is to find these local index sets in such a way that the section is also continuous when passing from one stratum of the Grassmannian to another. This continuity statement is shown in [12, Theorem 4.19].

Using the index set I , we may also calculate the initial degenerations of all points in the tropical Grassmannian and deduce that their tropical multiplicity is one everywhere. We will explain a general argument showing the existence of a section in this case in Sect. 7.2.

In [14], Draisma and Postingshel give a different proof of Theorem 4.1. They work with the affine cone over the Grassmannian $\text{Gr}(2, n)$, define the section on a suitable subset, and use torus actions and tropicalized torus actions to move it around. Also in this approach a maximality statement such as Lemma 4.5 is used.

5 Skeleta of Semistable Pairs

In the next three sections, we give an overview of the results in [22]. In order to compare polyhedral substructures of Berkovich spaces with tropicalizations, we start by generalizing Berkovich’s notion of skeleta. Such skeleta are induced by the incidence complexes of the special fibers of suitable models. We extend this notion by adding a horizontal divisor on the model. For the rest of this paper, we fix an algebraically closed ground field K which is complete with respect to a non-archimedean, non-trivial absolute value.

5.1 Integral Affine Structures

For curves, metrics play an important role in the comparison results between tropical and analytic varieties, as we have seen in Sect. 3. The right way to generalize this to higher dimensions is to consider integral affine structures. Let M be a lattice in the finite-dimensional real vector space $M_{\mathbb{R}} = M \otimes_{\mathbb{Z}} \mathbb{R}$, and let $N = \text{Hom}(M, \mathbb{Z})$ its dual, which is a lattice in the dual space $N_{\mathbb{R}} = N \otimes_{\mathbb{Z}} \mathbb{R}$ of $M_{\mathbb{R}}$. We denote the associated pairing by $\langle \cdot, \cdot \rangle : M_{\mathbb{R}} \times N_{\mathbb{R}} \rightarrow \mathbb{R}$. Recall that $\Gamma = \log |K^{\times}|$.

An *integral Γ -affine polyhedron* in $N_{\mathbb{R}}$ is a subset of $N_{\mathbb{R}}$ of the form

$$\Delta = \{v \in N_{\mathbb{R}} \mid \langle u_i, v \rangle + \gamma_i \geq 0 \text{ for all } i = 1, \dots, r\}$$

for some $u_1, \dots, u_r \in M$ and $\gamma_1, \dots, \gamma_r \in \Gamma$. Any face of an integral Γ -affine polyhedron Δ is again integral Γ -affine. An *integral Γ -affine polyhedral complex* in $N_{\mathbb{R}}$ is a polyhedral complex whose faces are integral Γ -affine.

An *integral Γ -affine function* on $N_{\mathbb{R}}$ is a function from $N_{\mathbb{R}}$ to \mathbb{R} which is of the form

$$v \mapsto \langle u, v \rangle + \gamma$$

for some $u \in M$ and $\gamma \in \Gamma$. More generally, let M' be a second finitely generated free abelian group and let $N' = \text{Hom}(M', \mathbb{Z})$. An *integral Γ -affine map* from $N_{\mathbb{R}}$ to $N'_{\mathbb{R}}$ is a function of the form $F = \phi^* + v$, where $\phi : M' \rightarrow M$ is a homomorphism, $\phi^* : N_{\mathbb{R}} \rightarrow N'_{\mathbb{R}}$ is the dual homomorphism extended to $N_{\mathbb{R}}$, and $v \in N' \otimes_{\mathbb{Z}} \Gamma$. If $N' = M' = \mathbb{Z}^m$ and $F = (F_1, \dots, F_m) : N_{\mathbb{R}} \rightarrow \mathbb{R}^m$ is a function, then F is integral Γ -affine if and only if each coordinate $F_i : N_{\mathbb{R}} \rightarrow \mathbb{R}$ is integral Γ -affine.

An *integral Γ -affine map* from an integral Γ -affine polyhedron $\Delta \subset N_{\mathbb{R}}$ to $N'_{\mathbb{R}}$ is defined as the restriction to Δ of an integral Γ -affine map $N_{\mathbb{R}} \rightarrow N'_{\mathbb{R}}$. If $\Delta' \subset N'_{\mathbb{R}}$ is an integral Γ -affine polyhedron, then a function $F : \Delta \rightarrow \Delta'$ is *integral Γ -affine* if the composition $\Delta \rightarrow \Delta' \hookrightarrow N'_{\mathbb{R}}$ is integral Γ -affine. We say that an integral Γ -affine map $F : \Delta \rightarrow N'_{\mathbb{R}}$ is *unimodular* if F is injective and if the inverse map $F(\Delta) \rightarrow \Delta$ is integral Γ -affine. Note that $F(\Delta)$ is an integral Γ -affine polyhedron in $N'_{\mathbb{R}}$.

5.2 Semistable Pairs

Note that since our ground field K is algebraically closed, the ring of integers K° is a valuation ring which is not discrete and not noetherian. We begin by describing the building blocks of the polyhedral substructures of Berkovich spaces which lend themselves to comparison with tropicalizations.

Let $0 \leq r \leq d$ be natural numbers and consider the K° -scheme

$$\mathcal{S} = \text{Spec}(K^\circ[x_0, \dots, x_d]/(x_0 \dots x_r - \pi))$$

for some $\pi \in K^\circ$ satisfying $|\pi| < 1$. Then \mathcal{S} is a flat scheme over K° with smooth generic fiber. Its special fiber contains $r+1$ irreducible components whose incidence complex is an r -dimensional simplex. Now fix a natural number $s \geq 0$ such that $r + s \leq d$, and consider the principal Cartier divisor $H(s) = \text{div}(x_{r+1}) + \dots + \text{div}(x_{r+s})$ on \mathcal{S} .

As explained in Sect. 2.2, the K° -scheme \mathcal{S} gives rise to two analytic spaces in the following way. On the one hand, we have an associated admissible formal scheme

$$S = \text{Spf}(K^\circ\{x_0, \dots, x_d\}/(x_0 \dots x_r - \pi)) \tag{5}$$

which we get by completing \mathcal{S} with respect to a non-zero element in K of absolute value < 1 . Its analytic generic fiber is $\mathcal{M}(A_S)$ for the affinoid algebra

$$A_S = K\{x_0, \dots, x_d\}/(x_0 \dots x_r - \pi).$$

It is a subset of the analytification of the generic fiber $\mathcal{S}_K = \text{Spec}(K[x_0, \dots, x_d]/(x_0 \dots x_r - \pi))$ of \mathcal{S} .

Now we look at the tropicalization map on $\mathcal{M}(A_S)$ which only takes into account the first $r + s + 1$ coordinates, i.e., the map

$$\begin{aligned} \text{Val} : \mathcal{M}(A_S) \setminus \mathcal{H} &\rightarrow \mathbb{R}_{\geq 0}^{r+s+1} \\ p &\mapsto (-\log |x_0(p)|, \dots, -\log |x_{r+s}(p)|), \end{aligned}$$

where \mathcal{H} is the support of the Cartier divisor induced by $H(s)$. Its image is $\Delta(r, \pi) \times \mathbb{R}_{\geq 0}^s$, where $\Delta(r, \pi) = \{(v_0, \dots, v_r) \in \mathbb{R}_{\geq 0}^{r+1} : v_0 + \dots + v_r = -\log |\pi|\}$ is a simplex in \mathbb{R}^{r+1} .

We define a continuous section $\Delta(r, \pi) \times \mathbb{R}_{\geq 0}^s \rightarrow \mathcal{M}(A_S)$ as follows. Note that the projection $(x_0, x_1, \dots, x_d) \mapsto (x_1, \dots, x_d)$ induces an isomorphism from $\mathcal{M}(A_S)$ to the affinoid subdomain $B = \{p : -\log |x_1(p)| - \dots - \log |x_r(p)| \leq -\log |\pi|\}$ of $\mathcal{M}(K\{x_1, \dots, x_d\})$.

Similarly, the projection $(v_0, v_1, \dots, v_{r+s}) \mapsto (v_1, \dots, v_{r+s})$ induces a homeomorphism

$$\Delta(r, \pi) \times \mathbb{R}_{\geq 0}^s \rightarrow \Sigma = \{(v_1, \dots, v_{r+s}) \in \mathbb{R}_{\geq 0}^{r+s} : v_1 + \dots + v_r \leq -\log |\pi|\}$$

For every $v = (v_1, \dots, v_{r+s}) \in \Sigma$ there is a bounded multiplicative norm $\| \cdot \|_v$ on $K\{x_1, \dots, x_d\}$ which is defined as follows:

$$\| \sum_{I=(i_1, \dots, i_d)} a_I x^I \|_v = \max_I \{ |a_I| \exp(-i_1 v_1 - \dots - i_{r+s} v_{r+s}) \}.$$

It satisfies $-\log \|x_1\|_v - \dots - \log \|x_r\|_v = v_1 + \dots + v_r \leq -\log |\pi|$. Therefore the point $\| \cdot \|_v$ is contained in the affinoid domain B . Hence there is a uniquely determined continuous map

$$\sigma : \Delta(r, \pi) \times \mathbb{R}_{\geq 0}^s \rightarrow \mathcal{M}(A_S)$$

making the diagram

$$\begin{array}{ccc}
 \Delta(r, \pi) \times \mathbb{R}_{\geq 0}^s & \xrightarrow{\sigma} & \mathcal{M}(A_S) \\
 \downarrow & & \downarrow \\
 \Sigma & \xrightarrow{v \mapsto || \cdot ||_v} & B
 \end{array}$$

commutative. The map σ is by construction a section of the map Val . We define $S(\mathcal{S}, H(s)) \subset \mathcal{M}(A_S) \subset (\mathcal{S}_K)^{\text{an}}$ as the image of σ and call it the skeleton of the pair $(\mathcal{S}, H(s))$.

Now we consider schemes which étale locally look like some \mathcal{S} .

Definition 5.1. A *strictly semistable pair* (\mathcal{X}, H) consists of an irreducible proper flat scheme \mathcal{X} over the valuation ring K° and a sum $H = H_1 + \dots + H_S$ of effective Cartier divisors H_i on \mathcal{X} such that \mathcal{X} is covered by open subsets \mathcal{U} which admit an étale morphism

$$\psi : \mathcal{U} \longrightarrow \mathcal{S} = \text{Spec}(K^\circ[x_0, \dots, x_d]/(x_0 \cdots x_r - \pi)) \tag{6}$$

for some $r \leq d$ and $\pi \in K^\times$ with $|\pi| < 1$. We assume that each H_i has irreducible support and that the restriction of H_i to \mathcal{U} is either trivial or defined by $\psi^*(x_j)$ for some $j \in \{r + 1, \dots, d\}$.

This is a generalization of de Jong’s notion of a strictly semistable pair over a discrete valuation ring [13], if we add the divisor of the special fiber to the horizontal divisor H . It is sometimes convenient to include only the horizontal divisor H as part of the data, since it is a Cartier divisor, whereas the special fiber of \mathcal{X} may not be one. For a more detailed discussion of this issue see [22, Proposition 4.17]. If $d = 1$, i.e., in the case of curves, semistable models and therefore nice skeletons are always available after an extension of the ground field. In higher dimensions, we have to allow alterations in order to find semistable models, see [13, Theorem 6.5].

5.3 Skeletons

Let (\mathcal{X}, H) be a semistable pair in the sense of Definition 5.1. The generic fiber X of \mathcal{X} is smooth of dimension d . Hence d is constant in every chart \mathcal{U} , whereas the numbers r and s may vary with \mathcal{U} . We denote the special fiber of \mathcal{X} by $\mathcal{X}_s = \mathcal{X} \otimes_{K^\circ} \tilde{K}$.

Denote by V_1, \dots, V_R the irreducible components of the special fiber \mathcal{X}_s of \mathcal{X} . The Cartier divisors H_i which are part of our data give rise to horizontal closed subschemes \mathcal{H}_i of \mathcal{X} , which are locally cut out by a defining equation of H_i . Putting $D_i = V_i$ for $i = 1, \dots, R$ and $D_{i+R} = \mathcal{H}_i$, we get a Weil divisor $D = \sum_{i=1}^{R+S} D_i$ on \mathcal{X} . This gives rise to a stratification of \mathcal{X} , where a stratum is

defined as an irreducible component of a set of the form $\bigcap_{i \in I} D_i \setminus \bigcup_{i \notin I} D_i$ for some $I \subset \{1, \dots, R + S\}$. We call any stratum contained in the special fiber \mathcal{X}_s a vertical stratum and denote the set of all vertical strata by $\text{str}(\mathcal{X}_s, H)$.

Now we want to glue skeleta of local charts together in such a way that the faces of the resulting polyhedral complex are in bijective correspondence with the vertical strata in $\text{str}(\mathcal{X}_s, H)$. In order to achieve this, we may have to pass to a smaller covering in the category of admissible formal schemes.

Let (\mathcal{X}, H) be a strictly semistable pair with a covering as in Definition 5.1. We consider the induced formal open covering of the associated admissible formal scheme \mathcal{X} which is defined by completion. Hence \mathcal{X} is covered by formal open subsets \mathcal{U} which admit an étale morphism $\psi : \mathcal{U} \rightarrow \mathcal{S}$ to a formal scheme \mathcal{S} as in (5).

It is shown in [22, Proposition 4.1] that, after passing to a refinement, the formal étale covering

$$\psi : \mathcal{U} \rightarrow \mathcal{S}$$

has the property that $\psi^{-1}\{x_0 = \dots = x_{r+s} = 0\}$ is a vertical stratum S in the special fiber \mathcal{X}_s such that for every vertical stratum T the following condition holds: The closure \bar{T} of T in \mathcal{X}_s meets \mathcal{U}_s if and only if $S \subset \bar{T}$.

We define the skeleton of $(\mathcal{U}, H|_{\mathcal{U}})$ as the preimage of the skeleton of the standard pair: $S(\mathcal{U}, H|_{\mathcal{U}}) = \psi^{-1}(S(\mathcal{S}, H(s)))$. This is a subset of the analytic generic fiber \mathcal{U}_η of \mathcal{U} which does not meet the horizontal divisor H .

It follows from results of Berkovich [5] that the étale map ψ actually induces a homeomorphism between $S(\mathcal{U}, H|_{\mathcal{U}})$ and $S(\mathcal{S}, H(s)) \simeq \Delta(\pi, r) \times \mathbb{R}_{\geq 0}^s$. Hence the map

$$\text{Val} \circ \psi : S(\mathcal{U}, H|_{\mathcal{U}}) \rightarrow \Delta(\pi, r) \times \mathbb{R}_{\geq 0}^s$$

is a homeomorphism.

It is shown in [22, Sect. 4.5] that the skeleton $S(\mathcal{U}, H|_{\mathcal{U}})$ only depends on the minimal stratum S contained in the special fiber of \mathcal{U} . Therefore we denote it by Δ_S and call it the *canonical polyhedron of S* (see Fig. 4). The dimensions of S and of the canonical polyhedron Δ_S are defined in an obvious way and add up to d , see [22, Proposition 4.10].

As we have seen, Δ_S is homeomorphic to $\Delta(\pi, r) \times \mathbb{R}_{\geq 0}^s$ for suitable data r, s , and π as above. We call $\Delta(\pi, r)$ the *finite part* and $\mathbb{R}_{\geq 0}^s$ the *infinite part* of Δ_S .

As explained above, $S \mapsto \Delta_S$ is a bijective correspondence between faces Δ_S of the skeleton $S(\mathcal{X}, H)$ and vertical strata S induced by the divisor D in the special fiber \mathcal{X}_s . Recall that D is given by the horizontal divisor H plus all irreducible components of \mathcal{X} (which we also call vertical divisors). Note that a vertical stratum T satisfies $S \subset \bar{T}$ if and only if Δ_T is a closed face of Δ_S .




Fig. 4 The canonical polyhedron Δ_S in the case $r = 0$ and $s = 2$, in the case $r = s = 1$, and in the case $r = 2$ and $s = 0$ (from left to right)

Now we can glue Δ_S and $\Delta_{S'}$ along the union of all canonical polyhedra associated with vertical strata R such that $\bar{R} \supset S \cup S'$. In this way we define the skeleton of (\mathcal{X}, H) as the union of all Δ_S for strata S in the special fiber \mathcal{X}_s :

$$S(\mathcal{X}, H) = \bigcup_s \Delta_S.$$

We obtain a piecewise linear space $S(\mathcal{X}, H)$ whose charts are integral Γ -affine polyhedra.

The skeleton $S(\mathcal{X}, H)$ is a closed subset of the analytic space $X^{\text{an}} \setminus H_K$, where we write H_K for the generic fiber of the support of H on \mathcal{X} . Using methods from [5], one can prove that it is in fact a strong deformation retract of $X^{\text{an}} \setminus H_K$. In [22, Theorem 4.13] it is shown that the retraction map can be extended to a retraction map from X^{an} to a suitable compactification of the skeleton.

If the horizontal divisor $H = 0$, then the skeleton $S(\mathcal{X}, 0)$ is equal to Berkovich’s skeleton of a semistable scheme, see [5].

6 Functions on the Skeleton

Throughout this section we fix a strictly semistable pair (\mathcal{X}, H) . We use the notation from the previous sections.

We want to show that for every non-zero rational function f on X such that the support of the divisor of f is contained in H_K , the function $-\log |f|$ factors through a piecewise integral Γ -affine map on the skeleton. Moreover, we will show a slope formula for this map.

Theorem 6.1 ([22, Proposition 5.2]). *Let f be a non-zero rational function on X such that the support of $\text{div}(f)$ is contained in H_K . We put $U = X \setminus H_K$ and consider the function*

$$F = -\log |f| : U^{\text{an}} \rightarrow \mathbb{R}.$$

Then F factors through the retraction map $\tau : U^{\text{an}} \rightarrow S(\mathcal{X}, H)$ to the skeleton:

$$\begin{array}{ccc}
 U^{\text{an}} & \xrightarrow{F} & \mathbb{R} \\
 & \searrow \tau & \nearrow F|_{S(\mathcal{X}, H)} \\
 & & S(\mathcal{X}, H)
 \end{array}$$

Moreover, the restriction of F to $S(\mathcal{X}, H)$ is an integral Γ -affine function on each canonical polyhedron.

The theorem is proved by considering the formal building blocks with distinguished strata and using a result of Gubler [20, Proposition 2.11].

Recall that we denote the dimension of X by d . The slope formula for F is basically a balancing condition around each $(d - 1)$ -dimensional canonical polyhedron of the skeleton which involves slopes in the direction of all adjacent d -dimensional polyhedra.

Let Δ_S be a d -dimensional canonical polyhedron of the skeleton $S(\mathcal{X}, H)$ containing the $(d - 1)$ -dimensional canonical polyhedron Δ_T . Then the stratum S in the special fiber \mathcal{X}_s is contained in the closure \bar{T} (which is a curve), and it is obtained as a component of the intersection of \bar{T} with one additional irreducible component of the divisor D .

This component is either vertical, i.e., a component of the special fiber, or horizontal, i.e., given by some component of H . If it is vertical, then the finite part of $\Delta_S \simeq \Delta(r, \pi) \times \mathbb{R}_{\geq 0}^s$ (i.e., the simplex $\Delta(r, \pi)$) is strictly larger than the finite part of Δ_T . In this case we say that Δ_S extends Δ_T in a bounded direction. If the component is horizontal, then the infinite part of Δ_S (i.e., the product $\mathbb{R}_{\geq 0}^s$) is strictly larger than the infinite part of Δ_T . In this case we say that Δ_S extends Δ_T in an unbounded direction.

We define the *bounded degree* $\text{deg}_b(\Delta_T)$ as the number of canonical polyhedral Δ_S extending Δ_T in a bounded direction. Similarly, we define the *unbounded degree* $\text{deg}_u(\Delta_T)$ as the number of canonical polyhedra Δ_S extending Δ_T in an unbounded direction.

In the case $d = 1$, i.e., if X is a curve, all Δ_T are vertices. A one-dimensional canonical polyhedron Δ_S extending Δ_T in a bounded direction is simply an edge of length $v(\pi) > 0$. In this case, the stratum S is a component of the intersection of two irreducible components in the special fiber \mathcal{X}_s . A one-dimensional canonical polyhedron Δ_S extending Δ_T in an unbounded direction is a ray of the form $\mathbb{R}_{\geq 0}$. In this case, the stratum S is a component of the intersection of the irreducible component of the special fiber given by T with a horizontal component of the divisor. Hence the degree $\text{deg}_b(\Delta_T)$ is the number of bounded edges in Δ_T , and $\text{deg}_u(\Delta_T)$ is the number of unbounded rays starting in the vertex Δ_T .

Now we want to define multiplicities via intersection theory. Since \mathcal{X} is not noetherian, the standard intersection theory tools from algebraic geometry are not available. However, on admissible formal schemes one can use analytic geometry to associate Weil divisors to Cartier divisors, and there is a refined intersection product with Cartier divisors, see [18, 19] and Appendix A in [22].

Definition 6.2. For every vertex $u \in \Delta_T$ we denote by V_u the associated irreducible component of the special fiber \mathcal{X}_s . We define an integer $\alpha(u, \Delta_T)$ as follows:

If Δ_T has zero-dimensional finite part $\{u\}$, we simply put $\alpha(u, \Delta_T) = \text{deg}_b(\Delta_T)$.

If not, then T lies in at least two irreducible components of the special fiber \mathcal{X}_s , which means that the finite part of T is a simplex $\Delta(r, \pi)$ of dimension at least one. If \mathcal{U} is a suitable formal chart, it is shown in [22, Proposition 4.17] that there exists a unique effective Cartier divisor C_u on \mathcal{U} such that its Weil divisor is equal to $v(\pi)(V_u \cap \mathcal{U}_s)$. We define $-\alpha(u, \Delta_T)$ as the intersection number (C_u, \bar{T}) , which is equal to the degree of the pullback of the line bundle associated with C_u from \mathcal{U} to \bar{T} .

Note that Cartwright [9] introduced the intersection numbers $\alpha(u, \Delta_T)$ in the (noetherian) situation where K° is discrete valuation ring. He uses them to endow the compact skeleton $S(\mathcal{X})$ with the structure of a tropical complex, see [9, Definition 1.1]. This notion also involves a local Hodge condition which plays no role in our slope formula. The paper [9] develops a theory of divisors on tropical complexes and investigates their relation to algebraic divisors.

Here we also need multiplicities for rays in Δ_T , which are defined as one-dimensional faces of the unbounded part of T .

Definition 6.3. Let H_r be the horizontal component corresponding to the ray r in Δ_T . We put

$$\alpha(r, \Delta_T) = -(H_r, \bar{T}),$$

where we take the intersection product of \bar{T} with the Cartier divisor H_r on \mathcal{X} .

Let F be a function on the skeleton $S(\mathcal{X}, H)$, which is integral Γ -affine on each canonical polyhedron. Now we are ready to define outgoing slopes of F on a $(d-1)$ -dimensional canonical polyhedron Δ_T along a d -dimensional polyhedron Δ_S .

Definition 6.4. Let Δ_T be a $(d-1)$ -dimensional canonical polyhedron and let Δ_S be a d -dimensional canonical polyhedron of $S(\mathcal{X}, H)$ containing Δ_T . We denote by $\Delta(r, \pi)$ the finite part of Δ_S . Let $F : \Delta_S \rightarrow \mathbb{R}$ be an integral Γ -affine function.

- (i) If Δ_S extends Δ_T in a bounded direction, then there exists a unique vertex w of Δ_S not contained in Δ_T . We put

$$\text{slope}(F; \Delta_T, \Delta_S) = \frac{1}{v(\pi)} \left(F(w) - \frac{1}{\text{deg}_b(\Delta_T)} \sum_{u \in \Delta_T} \alpha(u, \Delta_T) F(u) \right),$$

where we sum over all vertices u in Δ_T .

- (ii) If Δ_S extends Δ_T in an unbounded direction, then there exists a unique ray s in Δ_S not contained in Δ_T , and we put

$$\text{slope}(F; \Delta_T, \Delta_S) = d_s F - \frac{1}{\text{deg}_u(\Delta_T)} \sum_{r \in \Delta_T} \alpha(r, \Delta_T) d_r F,$$

where we sum over all rays in Δ_T . For any ray r we denote by $d_r F$ the derivative of F along the primitive vector in the direction of r .

If X is a curve, $\Delta_T = u$ is a vertex and Δ_S is an edge with vertices u and w , then $\text{slope}(F; \Delta_T, \Delta_S) = \frac{1}{v(\pi)}(F(w) - F(u))$. If Δ_S is a ray s starting in u , then we simply have $\text{slope}(F; \Delta_T, \Delta_S) = d_s F$. In higher dimensions, the definition is more involved, since the naive slope $\frac{1}{v(\pi)}(F(w) - F(u))$ depends on the choice of a vertex u in Δ_T . Therefore we define a replacement for a weighted midpoint in Δ_T as $\frac{1}{\text{deg}_b(\Delta_T)} \sum_{u \in \Delta_T} \alpha(u, \Delta_T) u$. Note that this point does not necessarily lie in Δ_T , since $\alpha(u, \Delta_T)$ may be negative.

We can now formulate the slope formula for skeleta.

Theorem 6.5 ([22, Theorem 6.9]). *Let $f \in K(X)^\times$ be a non-zero rational function such that the support of $\text{div}(f)$ is contained in H_K . Let $F : S(\mathcal{X}, H) \rightarrow \mathbb{R}$ be the restriction of the function $-\log |f|$ to the skeleton. Then F is continuous and integral Γ -affine on each canonical polyhedron of $S(\mathcal{X}, H)$, and for all $(d - 1)$ -dimensional canonical polyhedra we have*

$$\sum_{\Delta_S \succ \Delta_T} \text{slope}(F; \Delta_T, \Delta_S) = 0,$$

where the sum runs over all d -dimensional canonical polyhedra Δ_S containing Δ_T .

If X is a curve, the slope formula basically says that the sum of all outgoing slopes along edges or rays in a fixed vertex is zero. In this case, the slope formula is shown in [3, Theorem 5.15]. It is a reformulation of the non-archimedean Poincaré–Lelong formula proven in Thuillier’s thesis [27, Proposition 3.3.15]. The Poincaré–Lelong formula is an equation of currents in the form $dd^c \log |f| = \delta_{\text{div}(f)}$, where dd^c is a certain distribution-valued operator. This version of the slope formula was generalized to higher dimensions in the ground-breaking paper [10], where a theory of differential forms and currents on Berkovich spaces is developed. The approach of [10] uses tropical charts and does not rely on models or skeleta. In higher dimensions we see no direct relation to our slope formula.

7 Faithful Tropicalizations

We will now investigate the relation between skeleta, which are polyhedral substructures of analytic varieties, and tropicalizations, which are polyhedral images of algebraic or analytic varieties.

7.1 Finding a Faithful Tropicalization for a Skeleton

We start with a strictly semistable pair (\mathcal{X}, H) with skeleton $S(\mathcal{X}, H)$ and generic fiber X . We consider rational maps $f : X \dashrightarrow \mathbb{G}_{m,K}^n$ from X to a split torus. If $U \subset X$ is a Zariski open subvariety where f is defined, then $f|_U : U \rightarrow \mathbb{G}_{m,K}^n$ induces a tropicalization $\text{Trop}_f(U)$ of U , which is defined as the image of the map $\text{trop} \circ f^{\text{an}} : U^{\text{an}} \rightarrow \mathbb{R}^n$ as in Sect. 2.3.

Note that the skeleton is contained in the analytification of every Zariski open subset of X , since it only contains norms on the function field of X .

Definition 7.1. A rational map $f : X \dashrightarrow \mathbb{G}_{m,K}^n$ from X to a split torus is called a faithful tropicalization of the skeleton $S(\mathcal{X}, H)$ if the following conditions hold:

- (i) The map $\text{trop} \circ f^{\text{an}}$ is injective on $S(\mathcal{X}, H)$.
- (ii) Each canonical polyhedron Δ_S of $S(\mathcal{X}, H)$ can be covered by finitely many integral Γ -affine polyhedra such that the restriction of $\text{trop} \circ f^{\text{an}}$ to each of those polyhedra is a unimodular integral Γ -affine map.

For the definition of unimodular integral affine maps see Sect. 5.1.

It is easy to see that a rational map which is unimodular on the skeleton stays unimodular if we enlarge it with more rational functions on X , see [22, Lemma 9.3].

If we look at a building block $\mathcal{S} = \text{Spec}K^{\circ}[x_0, \dots, x_d]/(x_1 \dots x_r - \pi)$ as in Definition 5.1 with the local tropicalization $\Delta(r, \pi) \times \mathbb{R}_{\geq 0}^s$, we find that the coordinate functions x_0, \dots, x_{r+s} induce a faithful tropicalization. Collecting the corresponding rational functions on X for all canonical polyhedra of the skeleton, we get a rational map on X which is locally unimodular on the skeleton. It is shown in the proof of Theorem 9.5 of [22] how to enlarge this collection of rational function in order to ensure injectivity of the tropicalization map on the skeleton. In this way one can show the following result.

Theorem 7.2 ([22, Theorem 9.5]). *Let (\mathcal{X}, H) be a strictly semistable pair. Then there exists a collection of non-zero rational functions f_1, \dots, f_n on X such that the resulting rational map $f = (f_1, \dots, f_n) : X \dashrightarrow \mathbb{G}_{m,K}^n$ is a faithful tropicalization of the skeleton $S(\mathcal{X}, H)$.*

7.2 Finding a Copy of the Tropicalization Inside the Analytic Space

Let us now start with a given tropicalization of a very affine K -variety U , i.e., with a closed immersion $\varphi : U \hookrightarrow \mathbb{G}_{m,K}^n$ of a variety U in a split torus. As in Sect. 2.3 we consider the tropicalization map

$$\text{trop}_{\varphi} = \text{trop} \circ \varphi^{\text{an}} : U^{\text{an}} \rightarrow \mathbb{R}^n.$$

Recall that for every point $\omega \in \text{Trop}_\varphi(U)$ the preimage $\text{trop}_\varphi^{-1}(\omega) \subset U^{\text{an}}$ of the tropicalization map is the Berkovich spectrum of an affinoid algebra A_ω .

Lemma 7.3 ([22, Lemma 10.3]). *If the tropical multiplicity of ω is equal to one, then A_ω contains a unique Shilov boundary point.*

This lemma follows basically from [2], Remark after Proposition 4.17, which shows how the preimage of tropicalization is related to initial degenerations.

Now we can show that on the locus of tropical multiplicity one there exists a natural section of the tropicalization map.

Theorem 7.4 ([22, Theorem 10.7]). *Let $Z \subset \text{Trop}_\varphi(U)$ be a subset such that the tropical multiplicity of every point in Z is equal to one. By Lemma 7.3 this implies that for every $\omega \in Z$ the affinoid space $\text{trop}_\varphi^{-1}(\omega)$ has a unique Shilov boundary point which we denote by $s(\omega)$.*

The map $s : Z \rightarrow U^{\text{an}}$, given by $\omega \rightarrow s(\omega)$, is continuous and a partial section of the tropicalization map, i.e., on Z we have $\text{trop}_\varphi \circ s = \text{id}_Z$. Hence the image $s(Z)$ is a subset of U^{an} which is homeomorphic to Z .

Moreover, if Z is contained in the closure of its interior in $\text{Trop}_\varphi(U)$, then s is the unique continuous section of the tropicalization map on Z .

We give a sketch of the proof. It is enough to show that the section s is continuous and uniquely determined under our additional assumption. Since everything behaves nicely under base change we may assume that the valuation map $K^\times \rightarrow \mathbb{R}_{>0}$ is surjective. Let us first consider the case that $U = \mathbb{G}_{m,K}^n$ and φ the identity map. Then the section s is the identification of the tropicalization \mathbb{R}^n (which has multiplicity one everywhere) with the skeleton of $\mathbb{G}_{m,K}^n$ as defined in (3). It is clear from the explicit description of s in (3) that it is continuous in this case.

Moreover, if s' is a different section of the tropicalization map in the case $U = \mathbb{G}_{m,K}^n$ which satisfies $s(\omega) \neq s'(\omega)$, we find a Laurent polynomial f on which those two seminorms differ. Since $s(\omega)$ is the unique Shilov boundary point in the fiber of tropicalization, $s'(\omega)$ applied to f is strictly smaller than $s(\omega)$ applied to f . Since $s(\omega)$ is equal to $s'(\omega)$ on all monomials, it follows that the initial degeneration of f at ω cannot be a monomial. Therefore ω is contained in the tropical hypersurface $\text{Trop}(f)$. A continuity argument shows that the same argument works in a small neighborhood of ω . This is a contradiction since $\text{Trop}(f)$ has codimension one. Therefore s is indeed uniquely determined if φ is the identity map.

For general $\varphi : U \hookrightarrow \mathbb{G}_{m,K}^n$, where U has dimension $d \leq n$, one shows that there exists a linear map $\mathbb{Z}^n \rightarrow \mathbb{Z}^d$ with the following property: Let $\alpha : \mathbb{G}_{m,K}^n \rightarrow \mathbb{G}_{m,K}^d$ be the corresponding homomorphism of tori and consider $\psi = \alpha \circ \varphi : U \rightarrow \mathbb{G}_{m,K}^d$. Let $S(\mathbb{G}_{m,K}^d)$ be the skeleton of the torus as in (3). Then for all $\omega \in Z$ we have $\{s(\omega)\} = \text{trop}_\varphi^{-1}(\omega) \cap \psi^{-1}(S(\mathbb{G}_{m,K}^d))$. This implies that $s(Z) = \text{trop}_\varphi^{-1}(Z) \cap \psi^{-1}(S(\mathbb{G}_{m,K}^d))$ is closed, from which we can deduce continuity of s . Uniqueness follows from uniqueness in the torus case by composing the sections with ψ .

Note that the preceding theorem does not make any assumption on the existence of specific models. If we assume that (\mathcal{X}, H) is a semistable pair and $U = X \setminus H_K$, then the image of the section s defined in Theorem 7.4 is contained in the skeleton $S(\mathcal{X}, H)$. This is shown in [22, Proposition 10.9].

The preceding theorem treats the case of tropicalizations in tori. Let $\varphi : X \hookrightarrow Y$ be a closed embedding of X in a toric variety Y associated with the fan Δ . Then we may consider the associated tropicalization $\text{Trop}_\varphi(X)$ of X , i.e., the image of $\text{trop} \circ \varphi^{\text{an}} : X^{\text{an}} \rightarrow Y^{\text{an}} \rightarrow N_{\mathbb{R}}^{\Delta}$, where $N_{\mathbb{R}}^{\Delta}$ is the associated partial compactification of $N_{\mathbb{R}}$, see Sect. 2.3. We can apply Theorem 7.4 to all torus orbits. In this way, we get a section of the tropicalization map on the locus $Z \subset \text{Trop}_\varphi(X)$ of tropical multiplicity one which is continuous on the intersection with each toric stratum. It is a natural question under which conditions this section is continuous on the whole of Z . This might shed new light on Theorem 4.1 for the tropical Grassmannian.

Acknowledgements The author is very grateful to Walter Gubler for his helpful comments.

References

1. Baker, M., Rumely, R.: Potential theory and dynamics on the Berkovich projective line. In: *Mathematical Surveys and Monographs*, vol. 159. American Mathematical Society, Providence (2010)
2. Baker, M., Payne, S., Rabinoff, J.: Nonarchimedean geometry, tropicalization and metrics on curves. *Algebr. Geom.* **3**, 63–105 (2016)
3. Baker, M., Payne, S., Rabinoff, J.: On the structure of non-archimedean analytic curves. In: *Contemporary Mathematics*, vol. 605, pp. 93–121. American Mathematical Society, Providence (2013)
4. Berkovich, V.: Spectral theory and analytic geometry over non-Archimedean fields. In: *Mathematical Surveys and Monographs*, vol. 33. American Mathematical Society, Providence (1990)
5. Berkovich, V.: Smooth p -adic spaces are locally contractible. *Invent. Math.* **137**, 1–84 (1999)
6. Berkovich, V.: Smooth p -adic spaces are locally contractible. II. In: *Geometric Aspects of Dwork Theory*, pp. 293–370. De Gruyter, Berlin (2004)
7. Bosch, S.: *Lectures on Formal and Rigid Geometry*. Lecture Notes in Mathematics, vol. 2105. Springer, Heidelberg (2014)
8. Bosch, S., Güntzer, S., Remmert, R.: *Non-Archimedean Analysis*. Springer, Heidelberg (1984)
9. Cartwright, D.: Tropical Complexes. Preprint, <http://arxiv.org/abs/1308.3813> (2013)
10. Chambert-Loir, A., Ducros, A.: Formes différentielles réelles et courants sur les espaces de Berkovich. Preprint, <http://arxiv.org/abs/1204.6277> (2012)
11. Conrad, B.: Several approaches to non-archimedean geometry. In: *p -Adic Geometry*. University Lecture Series, vol. 45, pp. 9–63. American Mathematical Society, Providence (2008)
12. Cueto, M., Häbich, M., Werner, A.: Faithful tropicalization of the Grassmannian of planes. *Math. Ann.* **360**, 391–437 (2014)
13. de Jong, J.A.: Smoothness, semi-stability and alterations. *Publ. Math. IHES* **83**, 51–93 (1996)
14. Draisma, J., Postingshel, E.: Faithful tropicalization and torus actions. *Manuscripta Math.* **149**, 315–338 (2016)
15. Ducros, A.: Espaces analytiques p -adiques au sens de Berkovich. *Séminaire Bourbaki 2005/06*. Astérisque 311, exposé 958 (2007)
16. Einsiedler, M., Kapranov, M., Lind, D.: Non-Archimedean amoebas and tropical varieties. *J. Reine Angew. Math.* **601**, 139–157 (2006)
17. Foster, T., Gross, P., Payne, S.: Limits of tropicalizations. *Isr. J. Math.* **201**, 835–846 (2014)
18. Gubler, W.: Local heights of subvarieties over non-Archimedean fields. *J. Reine Angew. Math.* **498**, 61–113 (1998)

19. Gubler, W.: Local and canonical heights of subvarieties. *Ann. Scuola Norm. Super. Pisa Cl. Sci.* **2**(5), 711–760 (2003)
20. Gubler, W.: Tropical varieties for non-Archimedean analytic spaces. *Invent. Math.* **169**, 321–376 (2007)
21. Gubler, W.: A guide to tropicalizations. In: *Algebraic and Combinatorial Aspects of Tropical Geometry. Contemporary Mathematics*, vol. 589, pp. 125–189. American Mathematical Society, Providence (2013)
22. Gubler, W., Rabinoff, J., Werner, A.: Skeletons and tropicalizations. *Adv. Math.* **294**, 150–215 (2016)
23. Maclagan, D., Sturmfels, B.: *Introduction to Tropical Geometry*. American Mathematical Society, Providence (2015)
24. Payne, S.: Analytification is the limit of tropicalization. *Math. Res. Lett.* **16**, 543–556 (2009)
25. Speyer, D., Sturmfels, B.: The tropical Grassmannian. *Adv. Geom.* **4**, 389–411 (2004)
26. Temkin, M.: Introduction to Berkovich analytic spaces. In: *Berkovich Spaces and Applications. Lecture Notes in Mathematics*, vol. 2119, pp. 3–67. Springer, Heidelberg (2015)
27. Thuillier, A.: Théorie du potentiel sur les courbes on géométrie analytique non-archimédienne. Applications à la théorie d’Arakelov. Thèse Rennes. <http://tel.ccsd.cnrs.fr/documents/archives0/00/01/09/90/index.html> (2005)
28. Tyomkin, I.: Tropical geometry and correspondence theorem via toric stacks. *Math. Ann.* **353**, 945–995 (2012)

Berkovich Skeleta and Birational Geometry

Johannes Nicaise

Abstract We give a survey of joint work with Mircea Mustața and Chenyang Xu on the connections between the geometry of Berkovich spaces over the field of Laurent series and the birational geometry of one-parameter degenerations of smooth projective varieties. The central objects in our theory are the *weight function* and the *essential skeleton* of the degeneration. We tried to keep the text self-contained, so that it can serve as an introduction to Berkovich geometry for birational geometers.

Keywords Berkovich spaces • Skeleta • Weight function • Minimal model program

1 Introduction

Let R be a complete discrete valuation ring with residue field k and quotient field K . The main example to keep in mind is $R = \mathbb{C}[[t]]$. The discrete valuation on K gives rise to a non-archimedean absolute value on K that one can use to develop a theory of analytic geometry over K . The theory that we will use is the one introduced by Berkovich in [2]. The principal purpose of these notes is to describe some interactions between Berkovich geometry over K and the birational geometry of degenerations of algebraic varieties over R . For a nice introduction to related results over trivially valued base fields, we refer to [15].

In fact, we will use only a small part of the theory of Berkovich spaces: we are mainly interested in the underlying topological space of the analytification of an algebraic K -variety. The structure of this space can be described in terms of classical valuation theory; this will be explained in Sect. 2.

Let X be a connected, smooth, and proper K -variety of dimension n . We denote by X^{an} the Berkovich analytification of X . An *sncd*-model of X is a regular scheme \mathcal{X} of finite type over R , endowed with an isomorphism of K -schemes $\mathcal{X}_K \rightarrow X$, such that the special fiber \mathcal{X}_k is a divisor with strict normal crossings. To any

J. Nicaise (✉)

Department of Mathematics, KU Leuven, Celestijnenlaan 200B, 3001 Heverlee, Belgium
e-mail: johannes.nicaise@wis.kuleuven.be

proper *sncd*-model \mathcal{X} of X over R one can attach a subspace $\mathrm{Sk}(\mathcal{X})$ of X^{an} , called the Berkovich skeleton of X , which is canonically homeomorphic to the dual intersection complex of the strict normal crossings divisor \mathcal{X}_k . This skeleton can be viewed as the space of real valuations on the function field of X that extend the discrete valuation on K and that are monomial with respect to \mathcal{X}_k . The most important property of the skeleton $\mathrm{Sk}(\mathcal{X})$ is that it controls the homotopy type of X^{an} : it is a strong deformation retract of X^{an} . This provides an interesting link between the geometry of X^{an} and the birational geometry of models of X .

If X has dimension one and genus at least one, then X has a unique minimal *sncd*-model, and thus a canonical Berkovich skeleton. In higher dimensions, minimal *sncd*-models no longer exist. Nevertheless, one can ask whether it is still possible to construct a canonical skeleton in X^{an} . We will present two constructions, which we developed in collaboration with Mircea Mustaă [12] and Chenyang Xu [13], respectively. The first construction is based on work of Kontsevich and Soibelman on degenerations of Calabi–Yau varieties and mirror symmetry [10]; the second one relies on the Minimal Model Program, and in particular on the results in [6]. As we will see, both approaches yield the same result. We assume in the remainder of this introduction that the residue field k has characteristic zero.

The main idea behind the first approach is the following. Each proper *sncd*-model \mathcal{X} of X gives rise to a Berkovich skeleton $\mathrm{Sk}(\mathcal{X})$ in X^{an} . We will use pluricanonical forms ω on X to single out certain essential faces of the simplicial complex $\mathrm{Sk}(\mathcal{X})$, which must be contained in the skeleton of *every* proper *sncd*-model of X . The union of these ω -essential faces is called the Kontsevich–Soibelman skeleton of (X, ω) and denoted by $\mathrm{Sk}(X, \omega)$. It only depends on X and ω , but not on the choice of \mathcal{X} . Taking the union of the skeleta $\mathrm{Sk}(X, \omega)$ over all non-zero pluricanonical forms ω on X , we obtain a subspace of X^{an} with piecewise affine structure that we call the *essential skeleton* of X and that we denote by $\mathrm{Sk}(X)$.

This construction has the merit of being quite natural and elementary, but it is not at all clear from the definition what $\mathrm{Sk}(X)$ looks like or whether $\mathrm{Sk}(X)$ is still a strong deformation retract of X^{an} . Therefore, we will also consider a second approach. As we have already mentioned, minimal *sncd*-models usually do not exist if the dimension of X is at least two, but we can enlarge our class of models in such a way that minimal models exist and such that we can still use the members of this class to describe the homotopy type of X^{an} . The Minimal Model Program suggests to consider the so-called *dlt*-models of X , which should be viewed as proper *sncd*-models with mild singularities. We can define the skeleton $\mathrm{Sk}(\mathcal{X})$ of such a *dlt*-model by simply ignoring the singularities. The theory of minimal models guarantees that minimal *dlt*-models exist if the canonical sheaf of X is semi-ample, which means that some tensor power is generated by global sections. Minimal *dlt*-models are not unique, but the skeleton $\mathrm{Sk}(\mathcal{X})$ does not depend on the choice of a minimal *dlt*-model \mathcal{X} . By a careful analysis of the steps in the Minimal Model Program, it was proven in [6] that the skeleton $\mathrm{Sk}(\mathcal{X})$ can be obtained from the skeleton of any proper *sncd*-model of X by a sequence of elementary collapses. This implies that $\mathrm{Sk}(\mathcal{X})$ is still a strong deformation retract of X^{an} .

One of the main results of [13] is that these two constructions yield the same result: if the canonical sheaf of X is semi-ample, then the essential skeleton $\text{Sk}(X)$ of X coincides with the skeleton of any minimal dlt -model of X . In particular, $\text{Sk}(X)$ is a strong deformation retract of X^{an} . The semi-ample condition can be understood as follows: it guarantees that X has enough pluricanonical forms to detect all the important pieces of the skeleton of a proper $sncd$ -model.

Notation We denote by R a complete discrete valuation ring with maximal ideal \mathfrak{m} , residue field k , and quotient field K . We denote by v_K the discrete valuation $K^* \rightarrow \mathbb{Z}$. We define an absolute value on K by setting $|x|_K = \exp(-v_K(x))$ for every element x of K^* . A variety over a field F is a separated F -scheme of finite type. If α and β are elements of \mathbb{R}^m for some positive integer m , then we denote by $\alpha \cdot \beta$ their scalar product $\sum_{i=1}^m \alpha_i \beta_i$.

2 The Berkovich Skeleton of an $sncd$ -Model

2.1 Birational Points

(2.1.1) Let X be a connected and smooth K -variety of dimension n . We denote by X^{an} the Berkovich analytification of X and by $i : X^{\text{an}} \rightarrow X$ the analytification morphism. We will mainly be interested in the underlying topological space of X^{an} , which is easy to describe. As a set, X^{an} consists of the couples $(x, |\cdot|)$ where x is a scheme-theoretic point of X and $|\cdot| : \kappa(x) \rightarrow \mathbb{R}$ is an absolute value on the residue field $\kappa(x)$ of X at x that extends the absolute value $|\cdot|_K$ on K . The analytification map $i : X^{\text{an}} \rightarrow X$ is simply the forgetful map that sends a couple $(x, |\cdot|)$ to x . The topology on X^{an} is the coarsest topology such that the following two properties are satisfied:

- (1) the topology on X^{an} is finer than the Zariski topology, that is, the map $i : X^{\text{an}} \rightarrow X$ is continuous, and
- (2) for every Zariski-open subset U of X and every regular function f on U , the map

$$|f| : i^{-1}(U) \rightarrow \mathbb{R}^+ : (x, |\cdot|) \mapsto |f(x)|$$

is continuous.

Note that the definition of $|f|$ makes sense because $f(x)$ is an element of the residue field $\kappa(x)$. We will often denote a point of X^{an} simply by x , leaving the absolute value $|\cdot|$ implicit in the notation. It is convenient in many situations to switch between the multiplicative and additive viewpoint: we will denote by v_x the real valuation

$$v_x : \kappa(x)^* \rightarrow \mathbb{R} : a \mapsto -\ln |a|.$$

As usual, we extend it to zero by setting $v_x(0) = +\infty$. The *residue field* of X^{an} at a point x is defined as the completion of the residue field $\kappa(i(x))$ of X at $i(x)$ with respect to the absolute value $|\cdot|$. It is a complete valued extension of the field K , which we denote by $\mathcal{H}(x)$. The valuation ring of $\mathcal{H}(x)$ will be denoted by $\mathcal{H}(x)^\circ$. If f is a rational function on X that is defined at $i(x)$, then we can think of $f(i(x)) \in \kappa(i(x))$ as an element of $\mathcal{H}(x)$, and we denote this element by $f(x)$.

(2.1.2) The topological space X^{an} is Hausdorff, and it is compact if and only if X is proper over K . If X is a curve, then there exists a simple classification of the points on X^{an} and one can draw a fairly explicit picture of X^{an} ; see, for instance, Sect. 5.1 in [1]. If the dimension of X is at least two, it is much more difficult to give a precise description of the whole space X^{an} . We will see in the following sections, however, that one can produce many interesting points in this space using the birational geometry of X , and that these points suffice to control the homotopy type of X^{an} .

(2.1.3) The set X^{bir} of birational points of X^{an} is defined as the inverse image under $i : X^{\text{an}} \rightarrow X$ of the generic point of X . In the additive notation, this is simply the set of real valuations on the function field $K(X)$ of X that extend the discrete valuation v_K on K . We endow X^{bir} with the topology induced by the Berkovich topology on X^{an} . We will see in (2.4.13) that the inclusion $X^{\text{bir}} \rightarrow X^{\text{an}}$ is a homotopy equivalence if k has characteristic zero. By its very definition, X^{bir} is a birational invariant of X , so we can hope to recover interesting birational invariants of X from this topological space.

2.2 Models

(2.2.1) We will define certain subclasses of birational points using the geometry of R -models of X . An R -model of X is a flat separated R -scheme of finite type \mathcal{X} endowed with an isomorphism of K -schemes $\mathcal{X}_K \rightarrow X$. Note that we do not impose any properness condition on X or \mathcal{X} . We say that \mathcal{X} is an *sncd*-model of X if \mathcal{X} is regular and its special fiber \mathcal{X}_k is a divisor with strict normal crossings.

(2.2.2) Let \mathcal{X} be an R -model of X and let x be a point of X^{an} . We say that x has a *center* on \mathcal{X} if the canonical morphism $\text{Spec } \mathcal{H}(x) \rightarrow X$ extends to a morphism $\text{Spec } \mathcal{H}(x)^\circ \rightarrow \mathcal{X}$. Such an extension is unique if it exists, by the valuative criterion of separatedness. If it exists, the center of x on \mathcal{X} is defined as the image of the closed point of $\text{Spec } \mathcal{H}(x)^\circ$ in \mathcal{X} and denoted by $\text{sp}_{\mathcal{X}}(x)$. Note that the point $\text{sp}_{\mathcal{X}}(x)$ always lies on the special fiber \mathcal{X}_k of \mathcal{X} because \mathfrak{m} is contained in the maximal ideal of $\mathcal{H}(x)^\circ$. We denote by $\widehat{\mathcal{X}}_\eta$ the set of points on X^{an} that have a center on \mathcal{X} . If \mathcal{X} is proper over R , then $\widehat{\mathcal{X}}_\eta = X^{\text{an}}$ by the valuative criterion of properness. The map

$$\text{sp}_{\mathcal{X}}: \widehat{\mathcal{X}}_\eta \rightarrow \mathcal{X}_k$$

is called the reduction map or specialization map. It has the peculiar property of being *anti-continuous*, which means that the inverse image of an open set is closed.

Example 2.2.3. If $X = \mathbb{A}_K^1 = \text{Spec } K[T]$ and $\mathcal{X} = \mathbb{A}_R^1$, then

$$\widehat{\mathcal{X}}_\eta = \{x \in X^{\text{an}} \mid |T(x)| \leq 1\}$$

and $\text{sp}_{\mathcal{X}}(x)$ is the reduction of $T(x) \in \mathcal{H}(x)^o$ modulo the maximal ideal of $\mathcal{H}(x)^o$ (viewed as a point of $\mathcal{X}_k = \text{Spec } k[T]$).

(2.2.4) Assuming a bit more technology [3, Sect. 1], we can give an equivalent description of $\widehat{\mathcal{X}}_\eta$ and $\text{sp}_{\mathcal{X}}$. If we denote by $\widehat{\mathcal{X}}$ the formal \mathfrak{m} -adic completion of \mathcal{X} , then the generic fiber of $\widehat{\mathcal{X}}$ is a compact analytic domain in X^{an} whose underlying set is precisely $\widehat{\mathcal{X}}_\eta$. The reduction map $\text{sp}_{\mathcal{X}}$ is the map underlying the specialization morphism of locally ringed spaces

$$\text{sp}_{\mathcal{X}}: \widehat{\mathcal{X}}_\eta \rightarrow \widehat{\mathcal{X}}.$$

2.3 Divisorial and Monomial Points

(2.3.1) Let \mathcal{X} be a normal R -model of X . If E is an irreducible component of the special fiber \mathcal{X}_k with generic point ξ , then the fiber $\text{sp}_{\mathcal{X}}^{-1}(\xi)$ consists of a unique point, which we call the divisorial point of X^{an} associated with (\mathcal{X}, E) . It is the birational point x on X^{an} that corresponds to the discrete valuation v_x on $K(X)$ with valuation ring $\mathcal{O}_{\mathcal{X}, \xi}$, normalized in such a way that v_x extends the discrete valuation v_K on K . Thus if f is a non-zero rational function on X , then

$$v_x(f) = \frac{1}{N} \text{ord}_E f$$

where N denotes the multiplicity of E in the Cartier divisor \mathcal{X}_k on \mathcal{X} and $\text{ord}_E f$ is the order of f along E . A point of X^{an} is called divisorial if it is the divisorial point associated with some couple (\mathcal{X}, E) as above. We will denote the set of divisorial points by $X^{\text{div}} \subset X^{\text{bir}}$. It is not difficult to show that the set X^{div} is dense in X^{an} ; see, for instance, [12, Proposition 2.4.9].

(2.3.2) The set of divisorial points is totally disconnected. We will define a more general class of points that should be viewed as some kind of interpolations between divisorial points: the monomial points. Let \mathcal{X} be an *sncd*-model of X and let E_1, \dots, E_r be distinct irreducible components of the special fiber \mathcal{X}_k with respective multiplicities N_1, \dots, N_r in \mathcal{X}_k , and assume that the intersection $\bigcap_{i=1}^r E_i$ is non-empty. Let $\alpha = (\alpha_1, \dots, \alpha_r)$ be a tuple of positive real numbers such that $\sum_{i=1}^r \alpha_i N_i = 1$ and let ξ be a generic point of $\bigcap_{i=1}^r E_i$ (by the definition of an *sncd*-model, this intersection is regular and of pure dimension $n + 1 - r$, but it is not necessarily connected).

Proposition 2.3.3. *There exists a unique minimal real valuation*

$$v: \mathcal{O}_{\mathcal{X},\xi} \setminus \{0\} \rightarrow \mathbb{R}^+$$

such that $v(T_i) = \alpha_i$ for every i in $\{1, \dots, r\}$ and every local equation $T_i = 0$ for E_i in \mathcal{X} at ξ .

(2.3.4) Proposition 2.3.3 can be proven by combining [12, Propositions 2.4.4 and 3.1.6]. We will not give a complete proof here, but only sketch how the valuation v can be constructed. For every i in $\{1, \dots, r\}$ we choose a local equation $T_i = 0$ of E_i in \mathcal{X} at ξ . Then the elements T_i form a regular system of local parameters in the local ring $\mathcal{O}_{\mathcal{X},\xi}$. It is not difficult to show that every element f in $\mathcal{O}_{\mathcal{X},\xi}$ can be written in the completed local ring $\widehat{\mathcal{O}}_{\mathcal{X},\xi}$ as a power series

$$f = \sum_{\beta \in \mathbb{N}^r} c_\beta T^\beta \tag{1}$$

where each coefficient c_β is either zero or a unit in $\widehat{\mathcal{O}}_{\mathcal{X},\xi}$. Such an expansion is not unique, but one can show that the expression

$$v(f) := \min\{\alpha \cdot \beta \mid \beta \in \mathbb{N}^r, c_\beta \neq 0\} \tag{2}$$

does not depend on any choices and that it defines a valuation v with the required properties. If R has equal characteristic, then the arguments can be simplified by using the fact that $\widehat{\mathcal{O}}_{\mathcal{X},\xi}$ is isomorphic to the power series ring $\kappa(\xi)[[T_1, \dots, T_r]]$ by Cohen’s structure theorem.

(2.3.7) The valuation v in Proposition 2.3.3 extends to a real valuation $v: K(X)^* \rightarrow \mathbb{R}$. It extends the discrete valuation v_K on K : if π is a uniformizer in R , then in the ring $\mathcal{O}_{\mathcal{X},\xi}$ we can write

$$\pi = u \prod_{i=1}^r T_i^{N_i}$$

with u a unit, so that $v(\pi) = \sum_{i=1}^r \alpha_i N_i = 1$. Thus v defines a birational point x on X^{an} , which we call the monomial point associated with the data

$$(\mathcal{X}, (E_1, \dots, E_r), \alpha, \xi). \tag{3}$$

The point x belongs to $\widehat{\mathcal{X}}_\eta$, and $\text{sp}_{\mathcal{X}}(x) = \xi$. We remark for later use that formula (2) can be generalized as follows: if

$$f = \sum_{\beta \in \mathbb{N}^r} c_\beta d_\beta T^\beta$$

where each coefficient c_β is either zero or a unit in $\widehat{\mathcal{O}}_{\mathcal{X}, \xi}$ and each coefficient d_β belongs to K , then

$$v(f) = \min\{v_K(d_\beta) + \alpha \cdot \beta \mid \beta \in \mathbb{N}^r, c_\beta \neq 0\} \tag{4}$$

since we can rewrite d_β as the product of

$$\pi^{v_K(d_\beta)} = (u \prod_{i=1}^r T_i^{N_i})^{v_K(d_\beta)}$$

with a unit in R to get an expansion for f of the form (1).

(2.3.10) A point on X^{an} is called monomial if it is the monomial point associated with a tuple of data as in (3); we will also say that such a point is monomial with respect to the model \mathcal{X} . If $r = 1$, then we get precisely the divisorial point associated with (\mathcal{X}, E_1) . Thus every divisorial point is monomial. Conversely, the monomial point associated with (3) is divisorial (possibly with respect to a different model \mathcal{X}') if and only if the parameters α_i all belong to \mathbb{Q} (see [12, Proposition 2.4.8]). The set of monomial points on X^{an} will be denoted by X^{mon} . We have the following inclusions:

$$X^{\text{div}} \subset X^{\text{mon}} \subset X^{\text{bir}} \subset X^{\text{an}}.$$

2.4 The Berkovich Skeleton

(2.4.1) Let \mathcal{X} be an *sncd*-model of X . We define the Berkovich skeleton of \mathcal{X} as the set of all points of X^{an} that are monomial with respect to \mathcal{X} , and we denote it by $\text{Sk}(\mathcal{X})$. By construction, the Berkovich skeleton $\text{Sk}(\mathcal{X})$ is a subspace of $\widehat{\mathcal{X}}_\eta \cap X^{\text{mon}}$. The importance of this object is that we can give an explicit description of the topology on $\text{Sk}(\mathcal{X})$ and that this suffices to understand the homotopy type of $\widehat{\mathcal{X}}_\eta$, as we will now explain.

(2.4.2) We first need to recall the definition of the *dual complex* of the strict normal crossings divisor \mathcal{X}_k . We write $\mathcal{X}_k = \sum_{i \in I} N_i E_i$ and for every non-empty subset J of I , we set $E_J = \cap_{i \in J} E_i$. The dual complex of \mathcal{X}_k is a simplicial complex¹ $|\Delta(\mathcal{X}_k)|$ whose simplices of dimension d correspond bijectively to the connected components of the regular k -varieties E_J where J runs through the set of subsets of I of cardinality $d + 1$. If J and J' are non-empty subsets of I , and C and C' are connected components of E_J and $E_{J'}$, respectively, then the simplex corresponding

¹To be precise, $|\Delta(\mathcal{X}_k)|$ is not a simplicial complex in the strict sense, because we allow, for instance, multiple edges between two vertices. This has no importance for the present exposition (and can always be remediated by blowing up \mathcal{X} at connected components of the subvarieties E_J , which gives rise to a stellar subdivision of the corresponding face of $|\Delta(\mathcal{X}_k)|$). In any case, $|\Delta(\mathcal{X}_k)|$ is the topological realization of a finite simplicial set.

to C is a face of the simplex corresponding to C' if and only if C contains C' . Thus the vertices of $|\Delta(\mathcal{X}_k)|$ correspond to the irreducible components E_i of \mathcal{X}_k , and we will denote the vertices accordingly by v_i , $i \in I$. If i and j are distinct elements of I , then the number of edges between v_i and v_j is the number of connected components of $E_i \cap E_j$, and so on. In this way, the dual complex $|\Delta(\mathcal{X}_k)|$ encodes the combinatorial structure of the intersections of prime components in \mathcal{X}_k . The dimension of $|\Delta(\mathcal{X}_k)|$ is at most n , the dimension of X . If X has dimension one, then the dual complex $|\Delta(\mathcal{X}_k)|$ is more commonly known as the *dual graph* of the special fiber \mathcal{X}_k .

Example 2.4.3. Assume that X has dimension one and that \mathcal{X}_k has four irreducible components E_1, E_2, E_3, E_4 such that E_1 intersects each of the other components in precisely one point and there are no other intersection points. Then $|\Delta(\mathcal{X}_k)|$ is a graph with four vertices v_1, v_2, v_3, v_4 with one edge between v_1 and v_i for $i = 2, 3, 4$ and no other edges.

If X has dimension two and \mathcal{X}_k is isomorphic to the union of the coordinate planes in \mathbb{A}_k^3 , then $|\Delta(\mathcal{X}_k)|$ is the standard 2-simplex.

(2.4.4) We will now construct a map

$$\Phi: |\Delta(\mathcal{X}_k)| \rightarrow \text{Sk}(\mathcal{X}).$$

For each $i \in I$, our map Φ sends the vertex v_i of $|\Delta(\mathcal{X}_k)|$ to the divisorial point associated with (\mathcal{X}, E_i) . In order to define Φ on the higher-dimensional faces of $|\Delta(\mathcal{X}_k)|$, we use monomial valuations to interpolate between these divisorial valuations, as follows. Let y be a point of $|\Delta(\mathcal{X}_k)|$. Then there exists a unique face τ of $|\Delta(\mathcal{X}_k)|$ such that y lies in the interior τ° of τ . By the construction of $|\Delta(\mathcal{X}_k)|$, the face τ corresponds to a connected component C of an intersection E_J for some subset J of I . We denote by ξ the generic point of C . The vertices of τ correspond precisely to the irreducible components E_i with $i \in J$. We can represent the point y by a tuple of barycentric coordinates $\beta \in \mathbb{R}^J$ where each coordinate β_i is a positive real number and their sum is equal to one. Now we define $\Phi(y)$ as the monomial point of X^{an} associated with the data

$$(\mathcal{X}, (E_i)_{i \in J}, (\beta_i/N_i)_{i \in J}, \xi).$$

(2.4.5) It is easy to see that Φ is a bijection, since we can give the following description of the map Φ^{-1} . Let x be a point of $\text{Sk}(\mathcal{X})$. Then $\text{sp}_{\mathcal{X}}(x)$ is a generic point ξ of E_J , for some uniquely determined non-empty subset J of I . For each $i \in J$ we choose a local equation $T_i = 0$ for E_i in \mathcal{X} at ξ , and we set $\alpha_i = v_x(T_i)$. Then $\Phi^{-1}(x)$ lies in the interior of the face τ of $|\Delta(\mathcal{X}_k)|$ corresponding to the connected component of ξ in E_J , and its tuple of barycentric coordinates is equal to $(\alpha_i N_i)_{i \in J}$. In fact, we can say more.

Proposition 2.4.6. *The map*

$$\Phi: |\Delta(\mathcal{X}_k)| \rightarrow \text{Sk}(\mathcal{X})$$

is a homeomorphism.

Proof. Since the source of Φ is compact and the target is Hausdorff, we only need to prove that Φ is continuous. This is not difficult: using the definition of the Berkovich topology in (2.1.1) and the explicit description of monomial valuations in (2.3.4), one immediately checks that Φ is continuous on the interior of each of the faces of $|\Delta(\mathcal{X}_k)|$. To get the continuity at the boundary faces, one uses the following easy observation. Suppose that some of the α_i are zero in the construction of the valuation v in (2.3.4), say, $\alpha_1, \dots, \alpha_s \neq 0$ and $\alpha_{s+1} = \dots = \alpha_r = 0$ for some $s < r$. Then the formula we gave in (2) defines the monomial valuation associated with \mathcal{X} , the components E_1, \dots, E_s , the parameters $\alpha_1, \dots, \alpha_s$, and the unique generic point of $E_1 \cap \dots \cap E_s$ whose closure contains ξ . For details, see [12, Section 2.4.6 and Proposition 3.1.4]. \square

(2.4.7) We can use the homeomorphism Φ to endow $\text{Sk}(\mathcal{X})$ with a piecewise \mathbb{Z} -affine structure; see [12, Sect. 3.2]. This structure can be defined intrinsically on X^{an} and is independent of the choice of the model \mathcal{X} . The induced piecewise \mathbb{Q} -affine structure is simply the one inherited from the faces of $|\Delta(\mathcal{X})|$. We will not use the finer \mathbb{Z} -affine structure so we will not recall its definition here. If f is a non-zero rational function on X , then the function

$$\text{Sk}(\mathcal{X}) \rightarrow \mathbb{R}: x \mapsto \ln |f(x)|$$

is continuous and piecewise affine.

(2.4.8) Proposition 2.4.6 gives an explicit description of the topological space $\text{Sk}(\mathcal{X})$. We will now explain how one can use this description to determine the homotopy type of $\widehat{\mathcal{X}}_\eta$. First, we construct a retraction

$$\rho_{\mathcal{X}}: \widehat{\mathcal{X}}_\eta \rightarrow \text{Sk}(\mathcal{X})$$

for the embedding of $\text{Sk}(\mathcal{X})$ in $\widehat{\mathcal{X}}_\eta$. Let x be a point of $\widehat{\mathcal{X}}_\eta$. Let J be the set of indices $i \in I$ such that E_i contains the center $\text{sp}_{\mathcal{X}}(x)$ of x on \mathcal{X} . We denote by C the connected component of x in E_J and by ξ the generic point of C . For each $i \in J$ we choose a local equation $T_i = 0$ for E_i in \mathcal{X} at $\text{sp}_{\mathcal{X}}(x)$, and we set $\alpha_i = v_x(T_i)$. Then $\rho_{\mathcal{X}}(x)$ is the monomial point in X^{an} associated with the data

$$(\mathcal{X}, (E_i)_{i \in J}, (\alpha_i)_{i \in J}, \xi).$$

In other words, it is the unique point of the skeleton such that the Zariski closure of its center contains the center of x and which gives the same valuation to each local defining equation of an irreducible component E_i of \mathcal{X}_k passing through $\text{sp}_{\mathcal{X}}(x)$.

It is an easy exercise to verify that $\rho_{\mathcal{X}}$ is continuous. The most fundamental result about Berkovich skeleta is the following theorem.

Theorem 2.4.9 (Berkovich, Thuillier). *There exists a continuous map*

$$H : \widehat{\mathcal{X}}_{\eta} \times [0, 1] \rightarrow \widehat{\mathcal{X}}_{\eta}$$

such that $H(\cdot, 0)$ is the identity, $H(\cdot, 1)$ is the map $\rho_{\mathcal{X}} : \widehat{\mathcal{X}}_{\eta} \rightarrow \text{Sk}(\mathcal{X})$, and $H(x, t) = x$ for every point x of $\text{Sk}(\mathcal{X})$ and every t in $[0, 1]$. Thus $\text{Sk}(\mathcal{X})$ is a strong deformation retract of $\widehat{\mathcal{X}}_{\eta}$.

Corollary 2.4.10. *If X is proper and \mathcal{X} is a proper *sncd*-model of X , then $\text{Sk}(\mathcal{X})$ is a strong deformation retract of X^{an} . In particular, X^{an} has the same homotopy type as the simplicial complex $|\Delta(\mathcal{X}_k)|$.*

Proof. This follows from the fact that $\widehat{\mathcal{X}}_{\eta} = X^{\text{an}}$ if \mathcal{X} is proper, and from Proposition 2.4.6. □

(2.4.11) Giving a proof of Theorem 2.4.9 goes beyond the scope of this survey, but we will work out an elementary example in Sect. 2.5. The origins of Theorem 2.4.9 are the results by Berkovich on skeleta of the so-called poly-stable formal schemes [4]. Berkovich used these skeleta to prove that smooth non-archimedean analytic spaces are locally contractible. An *sncd*-model \mathcal{X} (or rather, its formal m -adic completion) is not poly-stable unless the special fiber \mathcal{X}_k is reduced². Thus we cannot directly apply Berkovich’s result here. If R has equal characteristic and \mathcal{X} is defined over an algebraic curve, we explained in [13, Theorem 3.1.3] how one can deduce Theorem 2.4.9 from results by Thuillier on skeleta over trivially valued fields [17]. The general case can be proven by translating Thuillier’s toroidal methods into the language of log-geometry; details will be given in a forthcoming publication.

(2.4.12) If X is proper over K , then the existence of a proper *sncd*-model \mathcal{X} is known if k has characteristic zero (by Hironaka’s resolution of singularities), and also if k has arbitrary characteristic and X is a curve (by Lipman’s resolution of singularities for excellent schemes of dimension two). Most experts believe that it should exist in general, but at this moment, resolution of singularities in positive and mixed characteristic remains one of the big open problems in algebraic geometry. Corollary 2.4.10 implies in particular that the homotopy type of the dual complex $|\Delta(\mathcal{X}_k)|$ does not depend on the choice of the proper *sncd*-model \mathcal{X}_k . This is an analog of Thuillier’s generalization of Stepanov’s theorem in [17], saying that the homotopy type of the dual complex of a log resolution of a pair of algebraic varieties over a perfect field is independent of the choice of the log resolution.

(2.4.13) If \bar{X} is a smooth compactification of X and \mathcal{X} is a proper *sncd*-model of \bar{X} , then the explicit construction of the strong deformation retract H from

²However, the class of poly-stable formal schemes is much larger than the class of *sncd*-models with reduced special fiber.

Theorem 2.4.9 shows that it restricts to strong deformation retracts of X^{an} and X^{bir} onto $\text{Sk}(\mathcal{X})$ (in fact, $H(x, t)$ lies in X^{bir} for every x in \overline{X}^{an} and every $t > 0$). Thus the inclusions $X^{\text{bir}} \rightarrow X^{\text{an}}$ and $X^{\text{an}} \rightarrow \overline{X}^{\text{an}}$ are homotopy equivalences. In particular, the homotopy type of the analytification of a smooth K -variety is a birational invariant if k has characteristic zero.

2.5 The Deformation Retraction in a Basic Example

(2.5.1) We will give an explicit construction of the map H from Theorem 2.4.9 for the following elementary example:

$$\mathcal{X} = \text{Spec } R[T_1, T_2]/(T_1^{N_1} T_2^{N_2} - \pi),$$

with π a uniformizer in R and N_1 and N_2 positive integers. Then \mathcal{X} is an *sncd*-model for its generic fiber $X = \mathcal{X}_K$, and $\widehat{\mathcal{X}}_\eta$ is the set of points x in X^{an} such that $|T_1(x)| \leq 1$ and $|T_2(x)| \leq 1$. We denote by E_i the component of \mathcal{X}_k defined by $T_i = 0$ for $i = 1, 2$ and by O the unique intersection point of E_1 and E_2 . The dual complex $|\Delta(\mathcal{X}_k)|$ is the standard 1-simplex

$$\Delta_1 = \{(\lambda, 1 - \lambda) \in \mathbb{R}^2 \mid 0 \leq \lambda \leq 1\}$$

and the morphism Φ constructed in (2.4.4) sends $(1, 0)$ to the divisorial point associated with (\mathcal{X}, E_1) , $(0, 1)$ to the divisorial point associated with (\mathcal{X}, E_2) , and $(\lambda, 1 - \lambda)$ to the monomial point associated with

$$(\mathcal{X}, (E_1, E_2), \left(\frac{\lambda}{N_1}, \frac{1 - \lambda}{N_2}\right), O)$$

for all $\lambda \in]0, 1[$.

(2.5.2) The construction of the map H is best understood in terms of torus actions. We set $c = \text{gcd}(N_1, N_2)$ and $M_i = N_i/c$ for $i = 1, 2$, and we choose integers a_1 and a_2 such that $a_1 M_1 + a_2 M_2 = 1$. For every complete valued field extension $(L, |\cdot|_L)$ of K we set

$$\mathbb{G}_L = \{x \in (\text{Spec } L[U_1, U_2]/(U_1^{M_1} U_2^{M_2} - 1))^{\text{an}} \mid |U_1(x)| = 1\}$$

with the group structure given by componentwise multiplication. For every element t in the interval $[0, 1]$ we define a point $\gamma_L(t)$ in \mathbb{G}_L as follows. We will make use of the isomorphism of L -algebras

$$L[V, V^{-1}] \rightarrow L[U_1, U_2]/(U_1^{M_1} U_2^{M_2} - 1): V \mapsto U_1^{a_2} U_2^{-a_1}$$

whose inverse is given by $U_1 \mapsto V^{M_2}$ and $U_2 \mapsto V^{-M_1}$. We can write every polynomial f in $L[V]$ as a Taylor expansion

$$f = \sum_{i \geq 0} c_i (V - 1)^i$$

around the point 1, where the coefficients c_i lie in L . Then the point $\gamma_L(t)$ is fully determined by the property that

$$|f(\gamma_L(t))| = \max_{i \geq 0} |c_i|_L t^i.$$

In other words, the point $\gamma_L(t)$ is the sup-norm on the closed disc of radius t around the point 1 in a completed algebraic closure $\widehat{L^a}$ of L . Note that

$$|U_1(\gamma_L(t))| = |V^{M_2}(\gamma_L(t))| = 1$$

for every t so that $\gamma_L(t)$ is indeed a point of \mathbb{G}_L . The map

$$\gamma_L : [0, 1] \rightarrow \mathbb{G}_L : t \mapsto \gamma(t)$$

is a continuous path from $\gamma_L(0) = 1$ to $\gamma_L(1)$.

(2.5.3) The torus \mathbb{G}_K acts on $\widehat{\mathcal{X}}_\eta$ by componentwise multiplication, and we can use this action together with the paths γ_L to produce paths in $\widehat{\mathcal{X}}_\eta$. For every point x of $\widehat{\mathcal{X}}_\eta$, the action of \mathbb{G}_K gives rise to a continuous map

$$\mathbb{G}_{\mathcal{H}(x)} \rightarrow (\widehat{\mathcal{X}}_\eta) \times_K \mathcal{H}(x) : g \mapsto g \cdot x.$$

For every t in $[0, 1]$, we define $H(x, t)$ as the image of $\gamma_{\mathcal{H}(x)}(t) \cdot x$ under the projection map

$$(\widehat{\mathcal{X}}_\eta) \times_K \mathcal{H}(x) \rightarrow \widehat{\mathcal{X}}_\eta.$$

In this way, we obtain a map

$$H: \widehat{\mathcal{X}}_\eta \times [0, 1] \rightarrow \widehat{\mathcal{X}}_\eta : (x, t) \mapsto H(x, t).$$

The map H is continuous by continuity of the paths γ_L and of the torus action on $\widehat{\mathcal{X}}_\eta$.

(2.5.4) We can also give a more explicit and down-to-earth (but less conceptual) description of the map H . Let x be a point of $\widehat{\mathcal{X}}_\eta$ and let t be an element of $[0, 1]$. For notational convenience, we set $x_1 = T_1(x)$ and $x_2 = T_2(x)$; these are elements

of the residue field $\mathcal{H}(x)$ of X^{an} at x . Let f be an element of $K[T_1, T_2]$. Then we can write the Laurent polynomial $f(x_1V^{M_2}, x_2V^{-M_1})$ in $\mathcal{H}(x)[V, V^{-1}]$ as a rational function

$$f(x_1V^{M_2}, x_2V^{-M_1}) = \frac{1}{V^j} \sum_{i \geq 0} c_i(V - 1)^i,$$

where the coefficients c_i belong to the valued field $\mathcal{H}(x)$ and only finitely many of them are non-zero. The point $H(x, t)$ is fully characterized by the property

$$|f(H(x, t))| = \max_i |c_i|t^i.$$

We remark for later reference that

$$|f(H(x, t))| \geq |c_0| = |f(x)|. \tag{5}$$

(2.5.5) Now we prove that H is a strong deformation retract onto $\text{Sk}(\mathcal{X})$. Setting $t = 0$ we find $|f(H(x, 0))| = |f(x)|$ so that $H(x, 0) = x$. To compute $H(x, 1)$ we use the fact that for every complete valued extension $(L, |\cdot|_L)$ of K , the closed disc with radius one around 1 in \widehat{L}^a coincides with the closed disc with radius one around 0, so that

$$|g(\gamma_L(1))| = \max_i |c_i|_L$$

for every polynomial $g = \sum_{i \geq 0} c_iV^i$ in $L[V]$. In this way, we see that for every polynomial

$$f = \sum_{i,j \geq 0} c_{ij}T_1^i T_2^j$$

in $K[T_1, T_2]$, we have

$$|f(H(x, 1))| = \max_{i,j} |c_{ij}|_K |x_1|^i |x_2|^j,$$

or, in additive notation:

$$v_{H(x,1)}(f) = \min_{i,j} \{v_K(c_{ij}) + iv_x(T_1) + jv_x(T_2)\}.$$

Thus we find by using formula (4) that $H(x, 1)$ is the monomial point on X^{an} associated with

$$(\mathcal{X}, (E_1, E_2), (v_x(T_1), v_x(T_2)), \mathcal{O}).$$

This is precisely the image $\rho_{\mathcal{X}}(x)$ of x under the retraction $\rho_{\mathcal{X}}: \widehat{\mathcal{X}}_{\eta} \rightarrow \text{Sk}(\mathcal{X})$. Finally, we show that $H(x, t) = x$ for all t in $[0, 1]$ when x is a point of the skeleton $\text{Sk}(\mathcal{X})$, that is, a monomial point with respect to \mathcal{X} . Direct computation shows that $|T_1(H(x, t))| = |x_1|$ and $|T_2(H(x, t))| = |x_2|$. Combining the inequality (5) with the minimality property of monomial valuations in Proposition 2.3.3, we see at once that $H(x, t)$ must be equal to x .

3 Weight Functions and the Kontsevich–Soibelman Skeleton

3.1 The Work of Kontsevich and Soibelman

(3.1.1) In [10], Kontsevich and Soibelman proposed a new interpretation of mirror symmetry based on non-archimedean geometry over the field of complex Laurent series $\mathbb{C}((t))$. Their fundamental idea was to encode a part of the geometry of a one-parameter degeneration of complex Calabi–Yau varieties into a topological manifold endowed with a \mathbb{Z} -affine structure with singularities, and to interpret mirror symmetry as a certain combinatorial duality between such manifolds. They worked out in detail the case of degenerations of $K3$ -surfaces. Similar ideas were developed by Gross and Siebert in their theory of *toric degenerations*. Gross and Siebert replaced the use of non-archimedean geometry by methods from tropical and logarithmic geometry and extended the results for $K3$ -surfaces to higher-dimensional degenerations [7].

(3.1.2) An essential ingredient of the construction of Kontsevich and Soibelman is the following. We denote by Δ a small disc around the origin of the complex plane and we set $\Delta^* = \Delta \setminus \{0\}$. We denote by t a local coordinate on Δ centered at 0. Let X be a smooth projective family of varieties over Δ^* and let ω be a relative differential form of maximal degree on the family $X \rightarrow \Delta^*$. Kontsevich and Soibelman associated with these data a skeleton $\text{Sk}(X, \omega)$, which is a topological subspace of the Berkovich analytification of the $\mathbb{C}((t))$ -variety obtained from X by base change. If X is a family of Calabi–Yau varieties, then we set $\text{Sk}(X, \omega) = \text{Sk}(X)$ where ω is any relative volume form on X . This definition does not depend on the choice of ω .

(3.1.3) Kontsevich and Soibelman proved that $\text{Sk}(X, \omega)$ can be explicitly computed on any strict normal crossings model \mathcal{X} for X over Δ : it is a union of faces of the Berkovich skeleton $\text{Sk}(\mathcal{X})$ of the model \mathcal{X} on which ω is minimal in a suitable sense. Their proof relied on the Weak Factorization Theorem. It is interesting to note that, even though the Berkovich skeleton $\text{Sk}(\mathcal{X})$ from Sect. 2.4 heavily depends on the chosen model \mathcal{X} , the Kontsevich–Soibelman skeleton $\text{Sk}(X, \omega)$ only depends on X and ω . It singles out certain faces of $\text{Sk}(\mathcal{X})$ that must appear in the skeleton of *every* strict normal crossings model.

(3.1.4) In [12], Mircea Mustață and the author extended this construction to varieties over complete discretely valued fields K of arbitrary characteristic, and to pluricanonical forms ω . Our approach does not use the Weak Factorization Theorem but only relies on basic computations on valuations and canonical sheaves. Moreover, we proved that the skeleton of a Calabi–Yau variety over $\mathbb{C}((t))$ is always connected. An interesting gadget that appears in our work is the *weight function*

$$\text{wt}_\omega : X^{\text{an}} \rightarrow \mathbb{R} \cup \{+\infty\}$$

associated with a smooth and proper K -variety X and a pluricanonical form ω on X . This weight function is piecewise affine on the Berkovich skeleton of any strict normal crossings model of X and strictly increasing as one moves away from the Berkovich skeleton. The Kontsevich–Soibelman skeleton is precisely the set of points where wt_ω reaches its minimal value; see Sect. 3.3.

3.2 Log Discrepancies in Birational Geometry

(3.2.1) Our approach is inspired by interesting analogies with some fundamental invariants in birational geometry. Let X be a smooth complex variety and \mathcal{I} a coherent ideal sheaf on X . Let v be a divisorial valuation on X , that is, a positive real multiple of the discrete valuation ord_E on the function field $\mathbb{C}(X)$ associated with a prime divisor E on a normal birational modification Y of X . We denote by N the multiplicity of the scheme $Z(\mathcal{I}\mathcal{O}_Y)$ along E and by ν the multiplicity of E in the relative canonical divisor $K_{Y/X}$. We set

$$\text{wt}_{\mathcal{I}}(v) = \frac{\nu + 1}{N} \text{ if } N \neq 0 \text{ and } +\infty \text{ otherwise,}$$

and we call this value the weight of \mathcal{I} at v . Then the infimum of the values $\text{wt}_{\mathcal{I}}(v)$ at all divisorial valuations v on X is called the log-canonical threshold of the pair (X, \mathcal{I}) and denoted by $\text{lct}(X, \mathcal{I})$. This is a measure for the singularities of the zero locus $Z(\mathcal{I})$ of \mathcal{I} on X , and one of the most important invariants in birational geometry. We refer to [9] for more background.

(3.2.2) It is a fundamental fact that the log-canonical threshold of (X, \mathcal{I}) can be computed on a single log resolution of (X, \mathcal{I}) , i.e., a proper birational morphism $h : Y \rightarrow X$ such that Y is smooth, h is an isomorphism over the complement of $Z(\mathcal{I})$, and $Z(\mathcal{I}\mathcal{O}_Y)$ is a strict normal crossings divisor on Y . Namely, we have

$$\text{lct}(X, \mathcal{I}) = \min\{\text{wt}_{\mathcal{I}}(v)\}$$

where v runs over the divisorial valuations associated with the prime components of $Z(\mathcal{I}\mathcal{O}_Y)$. If this minimum is reached on a prime component E of $Z(\mathcal{I}\mathcal{O}_Y)$, then we say that E computes the log-canonical threshold of (X, \mathcal{I}) . If we denote by \mathcal{E} the union of such prime components E , then the Connectedness Theorem

of Shokurov and Kollár [9, Theorem 7.4] states that for every point x of $Z(\mathcal{I})$ and every sufficiently small open neighborhood U of x in $Z(\mathcal{I})$, the topological space $h^{-1}(U) \cap \mathcal{E}$ is connected. This was the main source of inspiration for our theorem on the connectedness of the skeleton of a Calabi–Yau variety over $\mathbb{C}((t))$ (Theorem 3.5.6).

(3.2.3) In [5] and [8], a function closely related to the weight function $\text{wt}_{\mathcal{I}}$ was extended from the set of divisorial valuations on X to the non-archimedean link of $Z(\mathcal{I})$ in X , that is, the analytic space over the field \mathbb{C} with the trivial absolute value that we obtain by removing the generic fiber of the \mathbb{C} -variety $Z(\mathcal{I})$ from the generic fiber of the formal completion of X along $Z(\mathcal{I})$. We have made a similar construction to define weight functions on analytic spaces over discretely valued fields; this construction will be explained in Sect. 3.4.

3.3 Definition of the Kontsevich–Soibelman Skeleton

(3.3.1) Let X be a connected, smooth, and proper K -variety of dimension n , and let ω be a non-zero m -pluricanonical form on X , that is, a non-zero element of $\omega_{X/K}^{\otimes m}(X)$. Let \mathcal{X} be a regular R -model of X , E an irreducible component of \mathcal{X}_k , and x the divisorial point on X^{an} associated with (\mathcal{X}, E) . The relative canonical sheaf $\omega_{\mathcal{X}/R}$ is a line bundle on \mathcal{X} that extends the canonical line bundle $\omega_{X/K}$ on X . The differential form ω on X defines a rational section of $\omega_{\mathcal{X}/R}^{\otimes m}$ and thus a divisor $\text{div}_{\mathcal{X}}(\omega)$ on \mathcal{X} . We denote by N the multiplicity of E in \mathcal{X}_k and by ν the multiplicity of E in $\text{div}_{\mathcal{X}}(\omega)$. We define the weight of ω at x by the formula

$$\text{wt}_{\omega}(x) = (\nu + m)/N.$$

This definition only depends on x and not on the choice of \mathcal{X} and E . We define the weight of X with respect to ω by

$$\text{wt}_{\omega}(X) = \inf\{\text{wt}_{\omega}(x) \mid x \in X^{\text{div}}\} \in \mathbb{R} \cup \{-\infty\}.$$

(3.3.2) A divisorial point x on X^{an} is called ω -essential if the weight function wt_{ω} reaches its minimal value at x , that is,

$$\text{wt}_{\omega}(X) = \text{wt}_{\omega}(x).$$

The skeleton $\text{Sk}(X, \omega)$ of the pair (X, ω) is defined as the closure of the set of ω -essential divisorial points in the space of birational points X^{bir} . It is obvious from the definition that $\text{Sk}(X, \omega)$ is a birational invariant of the pair (X, ω) , since the spaces X^{bir} and X^{div} and the weight function wt_{ω} are birational invariants.

3.4 Definition and Properties of the Weight Function

(3.4.1) Without suitable assumptions on the existence of resolutions of singularities, we cannot say much more about the skeleton $\text{Sk}(X, \omega)$; for instance, we cannot prove that $\text{Sk}(X, \omega)$ is non-empty. Therefore, we will assume from now on that k has characteristic zero or X is a curve. With the current state of affairs, these are the cases where resolution of singularities is known in the form that we need. In particular, it is known that every proper R -model of X can be dominated by a proper *sncd*-model of X .

(3.4.2) In Sect. 2.4 we have attached to each *sncd*-model \mathcal{X} of X its Berkovich skeleton $\text{Sk}(\mathcal{X})$. It was defined as the set of all birational points on X that are monomial with respect to the strict normal crossings divisor \mathcal{X}_k on \mathcal{X} . We have shown in Proposition 2.4.6 that the skeleton $\text{Sk}(\mathcal{X})$ is canonically homeomorphic to the dual complex $|\Delta(\mathcal{X}_k)|$ of the strict normal crossings divisor \mathcal{X}_k .

Theorem 3.4.3 (Proposition 4.5.5 in [12]). *There exists a unique smallest function*

$$\text{wt}_\omega : X^{\text{an}} \rightarrow \mathbb{R} \cup \{+\infty\}$$

with the following properties.

- (1) *The function wt_ω is lower semi-continuous.*
- (2) *Let \mathcal{X} be an *sncd*-model for X and let x be a point of the Berkovich skeleton $\text{Sk}(\mathcal{X})$. Let f be a rational function on \mathcal{X} such that, locally at $\text{sp}_{\mathcal{X}}(x)$, we have*

$$\text{div}(f) = \text{div}_{\mathcal{X}}(\omega) + m(\mathcal{X}_k)_{\text{red}}.$$

Then

$$\text{wt}_\omega(x) = -\ln |f(x)|.$$

In particular, wt_ω is continuous and piecewise affine on $\text{Sk}(\mathcal{X})$, and we get the same value as in (3.3.1) on divisorial points. Moreover, for all x in $\widehat{\mathcal{X}}_\eta$, we have

$$\text{wt}_\omega(x) \geq \text{wt}_\omega(\rho_{\mathcal{X}}(x))$$

with equality if and only if $x \in \text{Sk}(\mathcal{X})$.

- (3) *The restriction of wt_ω to X^{bir} is a birational invariant of (X, ω) .*

Proof. We only give a rough sketch of the arguments and refer to [12] for details. The formula in (2) can be used to extend the weight function wt_ω to the set X^{mon} of monomial points on X . Of course, each monomial point will belong to the Berkovich skeleta of several *sncd*-models, and one must show that the formula does not depend

on the choice of an *sncd*-model. Next, one proves that the inequality in (2) holds for monomial points. When x is any point of X^{an} , one sets

$$\text{wt}_\omega(x) = \sup_{\mathcal{X}} \{\text{wt}_\omega(\rho_{\mathcal{X}})\}$$

where \mathcal{X} runs through the set of proper *sncd*-models of X . Then one can prove that the resulting function wt_ω on X^{an} satisfies all the properties in the statement. \square

3.5 Computation of the Kontsevich–Soibelman Skeleton

(3.5.1) We can use the properties of the weight function in Theorem 3.4.3 to compute the Kontsevich–Soibelman skeleton $\text{Sk}(X, \omega)$ on a fixed proper *sncd*-model \mathcal{X} of X . The divisorial points are dense in each face of $\text{Sk}(\mathcal{X})$ (they are precisely the points with barycentric coordinates in \mathbb{Q}). Point (2) of the theorem immediately implies that $\text{Sk}(X, \omega)$ is the subspace of the compact space $\text{Sk}(\mathcal{X})$ consisting of the points where the continuous function $\text{wt}_\omega|_{\text{Sk}(\mathcal{X})}$ reaches its minimal value, because it says that the weight function is strictly increasing if we move away from the skeleton (recall that $\widehat{\mathcal{X}}_\eta = X^{\text{an}}$ if \mathcal{X} is proper over R). In particular, $\text{Sk}(X, \omega)$ is a non-empty compact topological space. We can make this description much more explicit, as follows.

(3.5.2) We write $\mathcal{X}_k = \sum_{i \in I} N_i E_i$. For each $i \in I$, we denote by v_i the multiplicity of E_i in the divisor $\text{div}_{\mathcal{X}}(\omega)$. Recall that each face of $\text{Sk}(\mathcal{X})$ corresponds to a connected component C of an intersection $E_J = \bigcap_{j \in J} E_j$ where J is a non-empty subset of I . We say that the face is ω -essential if

$$\frac{v_j + m}{N_j} = \min \left\{ \frac{v_i + m}{N_i} \mid i \in I \right\}$$

for every j in J and C is not contained in the Zariski closure in \mathcal{X} of the pluricanonical divisor $\text{div}_X(\omega)$ (the divisor of zeroes of ω on the K -variety X). The reason for the latter condition will be explained in the proof of Theorem 3.5.3.

Theorem 3.5.3 (Theorem 4.7.5 in [12]). *The weight of X with respect to ω is given by*

$$\text{wt}_\omega(X) = \min \left\{ \frac{v_i + m}{N_i} \mid i \in I \right\}$$

and the skeleton $\text{Sk}(X, \omega)$ is the union of the ω -essential faces of $\text{Sk}(\mathcal{X})$. In particular, this union only depends on X and ω , and not on the choice of the model \mathcal{X} .

Proof. This follows easily from the properties of the weight function described in Theorem 3.4.3(2). We have already explained in (3.5.1) that $\text{Sk}(X, \omega)$ is the locus of points in $\text{Sk}(\mathcal{X})$ where wt_ω reaches its minimal value. Recall that for every $i \in I$, the value of wt_ω at the vertex of $\text{Sk}(\mathcal{X})$ corresponding to E_i is given by $(v_i + m)/N_i$. The explicit formula for the weight function on $\text{Sk}(\mathcal{X})$ implies that it is piecewise affine and concave on every face of $\text{Sk}(\mathcal{X})$. It is affine on a face if and only if the corresponding subvariety C of \mathcal{X}_k is not contained in the closure of $\text{div}_X(\omega)$; see [12, Theorem 4.7.5] for details. Thus $\text{Sk}(X, \omega)$ is the union of the ω -essential faces of $\text{Sk}(\mathcal{X})$. \square

Example 3.5.4. Suppose that $R = \mathbb{C}[[t]]$ and that X is a $K3$ -surface over K . Assume that X has an *sncd*-model \mathcal{X} such that \mathcal{X}_k is reduced and $\omega_{\mathcal{X}/R}$ is trivial. Such models play an important role in the classification of semi-stable degenerations of $K3$ -surfaces by Kulikov [11] and Persson–Pinkham [16]. They have the special property that $\text{Sk}(X, \omega) = \text{Sk}(\mathcal{X})$ for every volume form ω on X , since all multiplicities N_i are equal to one, and all v_i are equal by the triviality of $\omega_{\mathcal{X}/R}$.

(3.5.5) In [12] we proved the following variant of the Shokurov–Kollár Connectedness Theorem. Like the proofs of Shokurov and Kollár, our proof is based on vanishing theorems: we proved generalizations of Kawamata–Viehweg Vanishing and Kollár’s Torsion-free Theorem for varieties over power series rings of characteristic zero in one variable by means of Greenberg approximation.

Theorem 3.5.6 (Corollary 5.5.4 in [12]). *Assume that the residue field k of K has characteristic zero. If X is a geometrically connected, smooth, and proper K -variety of geometric genus one, and ω is a non-zero canonical form on X , then $\text{Sk}(X, \omega)$ is connected.*

One can say much more using advanced tools from the Minimal Model Program, as we will explain in the following section.

4 The Essential Skeleton and the Minimal Model Program

4.1 The Essential Skeleton

(4.1.1) Throughout this section, we assume that the residue field k of K has characteristic zero. Let X be a connected, smooth, and projective K -variety (the projectivity condition is needed to apply results from the Minimal Model Program). Let \mathcal{X} be a proper *sncd*-model of X . We have seen in Theorem 3.5.3 that for every non-zero pluricanonical form ω on X , the Kontsevich–Soibelman skeleton

$\mathrm{Sk}(X, \omega)$ singles out certain faces of $\mathrm{Sk}(\mathcal{X})$ that do not depend on the model \mathcal{X} . In [12, Sect. 4.9] we defined the *essential skeleton* $\mathrm{Sk}(X)$ of X as the union of the Kontsevich–Soibelman skeleta $\mathrm{Sk}(X, \omega)$ for all non-zero pluricanonical forms ω on X .

(4.1.2) If $\omega_{X/K}$ is trivial and ω is a volume form on X , then it is not hard to see that $\mathrm{Sk}(X) = \mathrm{Sk}(X, \omega)$: multiplying ω with an element λ in K^* simply shifts the weight function by $v_K(\lambda)$, and for every $m > 0$, the space of m -pluricanonical forms on X is generated by the m th tensor power of ω . It is not true for general X , however, that $\mathrm{Sk}(X) = \mathrm{Sk}(X, \omega)$ for some fixed pluricanonical form ω on X .

(4.1.3) Without suitable conditions on X , we cannot hope that $\mathrm{Sk}(X)$ is a strong deformation retract of X^{an} . For instance, if X is rational (e.g., a projective space \mathbb{P}_K^n), then all pluricanonical forms on X are zero and the essential skeleton is empty. However, Chenyang Xu and the author proved in [13] that $\mathrm{Sk}(X)$ is a strong deformation retract of X^{an} if X has “enough” pluricanonical forms. Our proof is based on the Minimal Model Program, and in particular on the results in [6]. We will now briefly explain the main ideas.

4.2 *dlt-Models*

(4.2.1) If X is a curve of genus ≥ 1 , then X has a unique minimal *sncd*-model \mathcal{X} , and thus a canonical Berkovich skeleton $\mathrm{Sk}(\mathcal{X})$. If X has dimension at least two, however, minimal *sncd*-models no longer exist in general. In order to get a good notion of minimal model, we have to enlarge the class of *sncd*-models to the so-called (good) *dlt*-models. The abbreviation *dlt* stands for *divisorially log terminal*. The precise definition of a *dlt*-model \mathcal{X} is quite technical; we do not give it here but refer to [13, Sect. 2.3] instead. The basic idea is that we allow certain mild singularities on \mathcal{X} , in accordance with the general philosophy of the Minimal Model Program. The set of points of \mathcal{X} where \mathcal{X} is regular and \mathcal{X}_k is a strict normal crossings divisor is an open subscheme of \mathcal{X} that we denote by $\mathcal{X}^{\mathrm{sncd}}$. The definition of a *dlt*-model guarantees that $\mathcal{X}^{\mathrm{sncd}}$ is still sufficiently large to capture all the important information about the skeleton; we set

$$\mathrm{Sk}(\mathcal{X}) := \mathrm{Sk}(\mathcal{X}^{\mathrm{sncd}}) \subset X^{\mathrm{bir}}.$$

(4.2.2) A *dlt*-model \mathcal{X} of X is called *minimal* if the line bundle $\omega_{\mathcal{X}/R}((\mathcal{X}_k)_{\mathrm{red}})$ is semi-ample, which means that some power of this line bundle is generated by global sections. Fundamental theorems in birational geometry imply that X has a minimal *dlt*-model if and only if the canonical line bundle $\omega_{X/K}$ is semi-ample (we refer to [13, Theorem 2.3.6] for detailed references). To be precise, we should assume that X is defined over an algebraic curve because the necessary tools from

the Minimal Model Program have only been developed under that assumption, but we will ignore this issue here; if $\omega_{X/K}$ is trivial, one can get rid of the algebraicity condition by using tools from logarithmic geometry [13, Sect. 4.2].

(4.2.3) Minimal *dlt*-models are not unique, but they are closely related (birationally crepant) and the skeleton $\text{Sk}(\mathcal{X})$ does not depend on the choice of the minimal *dlt*-model \mathcal{X} . One of the main results in [13] is the following.

Theorem 4.2.4 (Theorem 3.3.3 in [13]). *If $\omega_{X/K}$ is semi-ample and \mathcal{X} is any minimal *dlt*-model of X , then the skeleton $\text{Sk}(\mathcal{X})$ is equal to the essential skeleton $\text{Sk}(X)$ of X .*

(4.2.5) If \mathcal{X}' is any proper *sncd*-model of X , then the Minimal Model Program tells us how to modify \mathcal{X}' into a minimal *dlt*-model \mathcal{X} by a series of divisorial contractions and flips. The effect of the steps in the Minimal Model Program on the Berkovich skeleton $\text{Sk}(\mathcal{X}')$ was carefully studied in [6], and these authors proved that $\text{Sk}(\mathcal{X})$ can be obtained from $\text{Sk}(\mathcal{X}')$ by means of a sequence of *elementary collapses*, combinatorial operations on simplicial complexes which are, in particular, strong deformation retracts. Since we already know that $\text{Sk}(\mathcal{X}')$ is a strong deformation retract of X^{an} by Theorem 2.4.9, we obtain the following result.

Theorem 4.2.6 (Corollary 3.3.5 in [13]). *If $\omega_{X/K}$ is semi-ample, then the essential skeleton $\text{Sk}(X)$ is a strong deformation retract of X^{an} .*

(4.2.7) Under certain conditions on X , one can use further results from the Minimal Model Program to obtain information about the topological properties of $\text{Sk}(X)$. For instance, if $\omega_{X/K}$ is trivial, the residue field k is algebraically closed, and the skeleton $\text{Sk}(X)$ has maximal dimension (that is, the same dimension as X), then results by Kollár and Kovács imply that $\text{Sk}(X)$ is a closed pseudo-manifold (see [13, Theorems 4.1.7 and 4.2.4]).

(4.2.8) It would be quite interesting to have a proof of Theorem 4.2.6 that does not make use of the Minimal Model Program and the arguments in [6], but instead uses the properties of the weight function from Sect. 3.4 and the geometric structure of the Berkovich space X^{an} . A possible approach is the following. Assume that $\omega_{X/K}$ is trivial and let ω be a volume form on X . Then the essential skeleton $\text{Sk}(X)$ of X is the locus where the weight function wt_ω on X^{an} reaches its minimal value, and we have seen in Theorem 3.4.3 that it is strictly increasing if one moves away from the Berkovich skeleton of any *sncd*-model of X . It is tempting to speculate that one can attach a gradient vector field on X^{an} to wt_ω that induces a flow that contracts X^{an} onto $\text{Sk}(X)$. We refer to [14] for partial results in this direction and applications to the study of motivic zeta functions.

Acknowledgements I am grateful to Sam Payne for helpful comments on an earlier version of this text.

References

1. Baker, M.: An introduction to Berkovich analytic spaces and non-Archimedean potential theory on curves. In: *p-Adic Geometry*. University Lecture Series, vol. 45, pp. 123–174. American Mathematical Society, Providence (2008)
2. Berkovich, V.G.: *Spectral Theory and Analytic Geometry over Non-archimedean Fields*. Mathematical Surveys and Monographs, vol. 33. American Mathematical Society, Providence (1990)
3. Berkovich, V.G.: Vanishing cycles for formal schemes. *Invent. Math.* **115**(3), 539–571 (1994)
4. Berkovich, V.G.: Smooth p -adic analytic spaces are locally contractible. *Invent. Math.* **137**(1), 1–84 (1999)
5. Boucksom, S., Favre, C., Jonsson, M.: Valuations and plurisubharmonic singularities. *Publ. Res. Inst. Math. Sci.* **44**(2), 449–494 (2008)
6. de Fernex, T., Kollár, J., Xu, C.: The dual complex of singularities. To appear in *Proceedings of the Conference in Honor of Yujiro Kawamata's 60th Birthday*. *Advanced Studies in Pure Mathematics*. arXiv:1212.1675
7. Gross, M.: Mirror symmetry and the Strominger-Yau-Zaslow conjecture. In: *Current developments in mathematics 2012*, pp. 133–191, International Press, Somerville (2013)
8. Jonsson, M., Mustață, M.: Valuations and asymptotic invariants for sequences of ideals. *Ann. Inst. Fourier.* **62**(6), 2145–2209 (2012)
9. Kollár, J.: Singularities of pairs. In: *Algebraic Geometry – Santa Cruz 1995*. *Proc. Sympos. Pure Math.*, vol. 62, Part 1, pp. 221–287. American Mathematical Society, Providence (1997)
10. Kontsevich, M., Soibelman, Y.: Affine structures and non-archimedean analytic spaces. In: Etingof, P., Retakh, V., Singer, I.M. (eds) *The Unity of Mathematics*. In Honor of the Ninetieth Birthday of I. M. Gelfand. *Progress in Mathematics*, vol. 244, pp. 312–385. Birkhäuser, Boston (2006)
11. Kulikov, V.: Degenerations of $K3$ surfaces and Enriques surfaces. *Izv. Akad. Nauk SSSR Ser. Mat.* **41**(5), 1008–1042 (1977)
12. Mustață, M., Nicaise, J.: Weight functions on non-archimedean analytic spaces and the Kontsevich-Soibelman skeleton. *Algebr. Geom.* **2**(3), 365–404 (2015)
13. Nicaise, J., Xu, C.: The essential skeleton of a degeneration of algebraic varieties. To appear in *Am. Math. J.* arXiv:1307.4041
14. Nicaise, J., Xu, C.: Poles of maximal order of motivic zeta functions. *Duke Math. J.* **165**(2), 217–243 (2016)
15. Payne, S.: Topology of nonarchimedean analytic spaces and relations to complex algebraic geometry. *Bull. Am. Math. Soc.* **52**(2), 223–247 (2015)
16. Persson, U., Pinkham, H.: Degeneration of surfaces with trivial canonical bundle. *Ann. Math.* **113**(1), 45–66 (1981)
17. Thuillier, A.: Géométrie toroïdale et géométrie analytique non archimédienne. Application au type d'homotopie de certains schémas formels. *Manuscr. Math.* **123**(4), 381–451 (2007)

Metrization of Differential Pluriforms on Berkovich Analytic Spaces

Michael Temkin

Abstract We introduce a general notion of a seminorm on sheaves of rings or modules and provide each sheaf of relative differential pluriforms on a Berkovich k -analytic space with a natural seminorm, called Kähler seminorm. If the residue field \bar{k} is of characteristic zero and X is a quasi-smooth k -analytic space, then we show that the maximality locus of any global pluricanonical form is a PL subspace of X contained in the skeleton of any semistable formal model of X . This extends a result of Mustață and Nicaise, because the Kähler seminorm on pluricanonical forms coincides with the weight norm defined by Mustață and Nicaise when k is discretely valued and of residue characteristic zero.

Keywords Berkovich analytic spaces • Kähler seminorms • Pluriforms • Skeletons

1 Introduction

1.1 Motivation

It often happens that a Berkovich space X possesses natural skeletons, which are, in particular, deformational retracts of X of finite topological type, see [4, Sect. 4.3] and [6]. In some special cases, such as the case of curves of positive genus or abelian varieties, there is a canonical (usually, minimal) skeleton. As a rule, skeletons are obtained from nice formal models, e.g. polystable or, more generally, log smooth ones, though we should mention for completeness that a different and very robust method of constructing skeletons was developed very recently by Hrushovski and Loeser, see [20].

In [25], Kontsevich and Soibelman constructed a canonical skeleton of analytic K3 surfaces over $k = \mathbf{C}((t))$ by use of a new method: the skeleton is detected as

M. Temkin (✉)

Einstein Institute of Mathematics, The Hebrew University of Jerusalem,
Giv'at Ram, Jerusalem 91904, Israel
e-mail: temkin@math.huji.ac.il

the extremality locus of the canonical form. In [28], this method was extended by Mustață and Nicaise as follows. If k is discretely valued and X is the analytification of a smooth and proper k -variety, they constructed norms on the pluricanonical sheaves $\omega_X^{\otimes n}$ and showed that the maximality locus of any non-zero pluricanonical form ϕ is contained in the skeleton associated with any semistable formal model of X . The union of the maximality loci of non-zero pluricanonical forms is called the essential skeleton of X in [28]. It is an important “combinatorial” subset of X , see [29], although it does not have to be a skeleton of X .

An advantage of the above approach is that it constructs a valuable combinatorial subset, the essential skeleton, in a canonical way. In particular, the only input is the metrization of pluricanonical sheaves, and no choice of a formal model is involved. Slightly ironically, one heavily exploits formal models to metrize $\omega_X^{\otimes n}$. On the one hand, existence of nice global formal models is not needed since it suffices for any so-called divisorial point x to find a sufficiently small domain that contains x and possesses a regular formal model. The latter problem is much easier and the construction works fine when $\text{char}(\tilde{k}) > 0$ and existence of nice global models is a dream. But on the other hand, this still leads to technical restrictions, including the assumption that k is discretely valued. In addition, the construction of the norm is not so geometric: first one defines it at divisorial points by use of formal models, and then extends it to the whole X by continuity.

The original aim of this project was to provide a natural local analytic construction of Mustață–Nicaise norm, which applies to all points on equal footing and eliminates various technical restrictions of their method. The basic idea is very simple: provide Ω_X with the maximal seminorm making the differential $d : \mathcal{O}_X \rightarrow \Omega_X$ a non-expansive map, and induce from it a seminorm on $\omega_X^{\otimes n} = (\wedge^d \Omega_X)^{\otimes n}$. Moreover, the same definition makes sense for any morphism $f : X \rightarrow S$, so the construction generalizes to the relative situation and no assumption on k and f is needed. This increases the flexibility even in the original setting; for example, one obtains a way to work with analytic families of proper smooth varieties. Unfortunately, implementation of the basic idea is not so simple due to lack of various foundations. So, a large part of this project is devoted to developing basic topics, including a theory of Kähler seminorms on modules of differentials, its application to real-valued fields, metrization of sheaves of modules, etc. On the positive side, we think that this foundational work will be useful for future research in non-archimedean geometry and related areas.

Once Kähler seminorms will have been defined, we will study the maximality locus of pluricanonical forms. Unfortunately, the assumption that $\text{char}(\tilde{k}) = 0$ seems unavoidable with current technique, but we manage to treat the non-discrete case as well. In particular, we only use a result à la de Jong (see [41, Theorem 3.4.1]) instead of the existence of semistable model (a result à la Hironaka). Finally, under the assumptions of [28], we compare the Kähler seminorm to the Mustață–Nicaise norm on the pluricanonical sheaves. Surprisingly, they coincide only when $\text{char}(\tilde{k}) = 0$, and in general they are related by a factor which up to a constant coincides with the log different of $\mathcal{H}(x)/k$.

1.2 Methods

At few places in the paper, including the definition of Kähler seminorms, we have to choose one method out of few possibilities. Let us discuss briefly what these choices are.

1.2.1 An Approach Via Unit Balls

One way to define a seminorm on a k° -vector space V is by using its unit ball V° , which is a k° -module. Technically, this leads to the easiest way to metrize a coherent \mathcal{O}_X -sheaf \mathcal{F} : just choose an \mathcal{O}_X° -submodule \mathcal{F}_X° (subject to simple restrictions). For example, if the valuation is not discrete, one can simply define the Kähler seminorm $\|\cdot\|_\Omega$ on $\Omega_{X/S}$ as the seminorm associated with the sheaf $\mathcal{O}_X^\circ d_{X/S}(\mathcal{O}_X^\circ)$, the minimal \mathcal{O}_X° -submodule of $\Omega_{X/S}$ containing $d_{X/S}\mathcal{O}_X^\circ$. Unfortunately, this definition is problematic when the valuation is discrete or trivial (though, see Remark 6.1.6). In order to consider ground fields with discrete or trivial valuations on the equal footing, we have to develop all basic constructions in terms of seminorms themselves. For example, $\|\cdot\|_\Omega$ can be characterized as the maximal seminorm such that the map $d : \mathcal{O}_X \rightarrow \Omega_{X/S}$ is non-expansive (see Lemma 6.1.4). Still, some implicit use of unit balls is made in Sect. 5, e.g. see Theorem 5.1.8.

1.2.2 Seminormed Algebras Versus Banach Algebras

Most seminormed rings in Berkovich geometry are Banach. Nevertheless, more general seminormed rings also show up, with the main example being the local rings $\mathcal{O}_{X,x}$ and their residue fields $\kappa(x)$. For this reason, it is technically much more convenient to work with seminormed rings and modules rather than their Banach analogues throughout the paper. In addition, it turns out to be important to consider only non-expansive homomorphisms, while classical Banach categories contain all bounded homomorphisms.

1.2.3 Metrization of Sheaves

There exist two ways to define seminorms on sheaves of rings or modules. In this paper we implement a sheaf theoretic approach, which applies to any site. We introduce the notions of (pre)sheaves of seminorms on a sheaf of abelian groups (resp. rings or modules) \mathcal{A} . In fact, this is equivalent to introducing (pre)sheaves of seminormed abelian groups with the underlying sheaf \mathcal{A} . Various operations, such as tensor products, are defined using sheafification. A slight technical complication of the method is that one has to consider unbounded seminorms.

A simpler ad hoc method to metrize sheaves on topological spaces or sites with enough points is to metrize the stalks in a semicontinuous way: for a section $s \in \mathcal{F}(U)$ let $|s| : U \rightarrow \mathbf{R}_{\geq 0}$ denote the function sending $x \in U$ to $|s|_x$, then all functions $|s|$ should be upper semicontinuous. All operations are then defined stalkwise. The main problem with applying this method to our case is that one has to work with all points of the G -topology site X_G , which is not a standard tool in Berkovich geometry. We describe these points in the end of the paper, but they are not used in our main constructions.

Remark 1.2.4. (i) It is more usual to consider only continuous metrics. For example, the definition of metrization in [9] requires that all sections $|s|$ are continuous. Nevertheless, it is the semicontinuity that encodes the condition that $| \cdot |$ is a sheaf, see Theorem 3.2.11.

(ii) In the case of Kähler seminorms, non-continuous functions $\|\phi\|_\Omega$ arise in the simplest cases. For example, already the function $\|dt\|_\Omega$ on the disc $\mathcal{M}(k\{t\})$ is not continuous; in fact, $\|dt\|_\Omega$ is the radius function, see Sect. 6.2.1.

1.3 Overview of the Paper and Main Results

In Sect. 2 we fix our notation and study seminorms on vector spaces over real-valued fields and modules over real valuation rings. Most of the material is probably known to experts but some of it is hard to find in the literature, especially the material on index of semilattices and content of torsion modules. Section 3 deals with metrization of sheaves. First, we study sheaves on arbitrary sites and then specialize to the case of G -sheaves on analytic spaces. In Sect. 4, we extend the theory of Kähler differentials to seminormed rings. In particular, given a non-expansive homomorphism of seminormed rings $A \rightarrow B$ we show that $d_{B/A} : B \rightarrow (\Omega_{B/A}, \| \cdot \|_\Omega)$ is the universal non-expansive A -derivation, where the Kähler seminorm $\| \cdot \|_\Omega$ is defined as the maximal seminorm making the homomorphism $d_{B/A}$ non-expansive.

In Sect. 5 we study the Kähler seminorm on a vector space $\Omega_{K/A}$, where K is a real-valued field. In Theorem 5.1.8 we show that $\Omega_{K^\circ/A^\circ}^{\log}$ modulo its torsion is an almost unit ball of the Kähler seminorm on $\Omega_{K/A}$. Thus, the study of Kähler seminorms is tightly related to study of $\Omega_{K^\circ/A^\circ}^{\log}$ and ramification theory. Although this is classical in the discretely valued case, [18, Chap. 6] is the only reference for such material when $|K^\times|$ is dense. Unfortunately, even loc.cit. does not cover all our needs, so we have to dig into the theory of real-valued fields, that makes this section the most technical in the paper. One of our main results there is that if K is dense in a real-valued field L then $\widehat{\Omega}_{K/A} \xrightarrow{\sim} \widehat{\Omega}_{L/A}$, see Theorem 5.6.6 and its corollary.

In Sect. 6 we define the Kähler seminorm $\| \cdot \|_\Omega$ on $\Omega_{X/S}$ for any morphism $f : X \rightarrow S$. This is done by sheafifying the presheaf of Kähler seminorms on affinoid domains and it follows easily from the definition that $\| \cdot \|_\Omega$ is the maximal seminorm on $\Omega_{X/S}$ making the map $d : \mathcal{O}_{X_G} \rightarrow \Omega_{X/S}$ non-expansive. In Theorem 6.1.8 we

show that the completed stalk $\widehat{\Omega}_{X/S,x}$ is isomorphic to $\widehat{\Omega}_{\mathcal{H}(x)/\mathcal{H}(s)}$, where $s = f(x)$. This provides the main tool for explicit work with the seminorm $\| \cdot \|_{\Omega}$ and its stalks. Note that the main ingredient in proving Theorem 6.1.8 is that $\widehat{\Omega}_{\kappa(x)/\kappa(s)} = \widehat{\Omega}_{\mathcal{H}(x)/\mathcal{H}(s)}$ by Theorem 5.6.6. Another fundamental property of Kähler seminorms is established in Theorem 6.1.13: Kähler seminorms are determined by the usual points of X (the G -analyticity condition from Sect. 3.3) and the functions $\|s\|_{\Omega}$ are semicontinuous with respect to the usual topology of X (and not only the G -topology). Some examples of Kähler seminorms are described in Sect. 6.2, including the ones demonstrating that Kähler seminorms behave weirdly when k possesses wildly ramified extensions. Also, we study the compatibility of $\| \cdot \|_{\Omega}$ with base change in Sect. 6.3. In particular, we show that it is compatible with restriction to the fibers and tame extensions of the ground field (Theorem 6.3.11 and its corollaries), but is incompatible with wild extensions of the ground field. Finally, we use $\| \cdot \|_{\Omega}$ to metrize the sheaves $S^m(\Omega_{X/S}^n)$ of relative pluriforms on X .

Section 7 is devoted to recalling basic facts about formal models and skeletons of Berkovich spaces. Then we study in Sect. 8 metrization of the sheaves $\omega_X^{\otimes m}$ of pluricanonical forms on a rig-smooth space X . In Corollary 8.1.3 we obtain a simple formula that evaluates Kähler seminorms at monomial points, and we deduce in Theorem 8.1.6 that the restrictions of geometric Kähler seminorms onto PL subspaces of X are PL. The main ingredient here is Theorem 9.4.8. In Theorem 8.3.3 we establish the connection between Kähler seminorm on pluricanonical sheaves and the weight norm of Mustařa and Nicaise. Finally, in Sect. 8.2 we study the maximality locus of a non-zero pluricanonical form with respect to the geometric Kähler seminorm. We prove that it is contained in the essential skeleton of X , see Theorem 8.2.4, and, if $\text{char}(\tilde{k}) = 0$, it is a PL subspace of X , see Theorem 8.2.9. When $\text{char}(\tilde{k}) = 0$ this extends the results of Mustařa and Nicaise to the case of non-discrete $|k^{\times}|$ and arbitrary quasi-smooth X , not necessarily algebraizable or even strictly analytic. (If $\text{char}(\tilde{k}) > 0$, our norm differs from the weight norm.)

Finally, Sect. 9 is devoted to study the topological realizations of the G -topologies on analytic spaces and PL spaces. Our description of the topological space $|X_G|$ seems to be new, see Sect. 9.1 and Remark 9.2.9(ii) concerning the connection to adic and reified adic spaces. In particular, we interpret the points of $|X_G|$ in terms of the graded reductions of germs. Also, we interpret points of the PL topologies in terms of combinatorial valuations (or valuations on lattices), see Sect. 9.3, and for a PL subspace $P \subseteq X$ we describe the embedding $|P_G| \hookrightarrow |X_G|$. This is used to prove a strong result on the structure of P : locally P possesses a residually unramified chart, see Theorem 9.4.8. It seems that our usage of the space $|P_G|$ is more or less equivalent to the use of model theory in [13, 16], see Remark 9.3.3. Note also that Theorem 9.4.8 is the only result of Sect. 9 used in the main part of the paper (in the proof of Theorem 8.1.6).

2 Real-Valued Fields

In this section we study seminormed vector spaces over a real-valued field K and modules over the ring of integers K° .

2.1 Conventions

First, let us fix basic terminology and notation on valued fields and analytic spaces.

2.1.1 Valued Fields

By a *valued field* we mean a field F provided with a non-archimedean valuation $|\cdot| : F \rightarrow \{0\} \cup \Gamma$, where Γ is an ordered group. The conditions $|x| \leq 1$ and $|x| < 1$ define the *ring of integers* F° and its maximal ideal $F^{\circ\circ}$, respectively. In addition, $\tilde{F} = F^\circ/F^{\circ\circ}$ denotes the *residue field* of F .

2.1.2 The Real-Valued Case

Assume that a valued field K is *real-valued*, i.e. $\Gamma = \mathbf{R}_{>0}$. Then $|\cdot|$ is a norm and hence defines a topology on K . We will use the notation $|K^{\circ\circ}| = \sup_{\pi \in K^{\circ\circ}} |\pi|$. Thus, $|K^{\circ\circ}| = 0$ if the valuation is trivial, $|K^{\circ\circ}| = |\pi|$ if K is discretely valued with uniformizer π , and $|K^{\circ\circ}| = 1$ otherwise.

Let, now, π be any element of $K^{\circ\circ} \setminus \{0\}$ if the valuation is non-trivial, and set $\pi = 0$ otherwise. In particular, the induced topology on K° is the π -adic one. Given a K° -module M we say that an element $x \in M$ is *divisible* if it is infinitely π -divisible. In particular, if the valuation of K is trivial, then 0 is the only divisible element. A K° -module is *divisible* if all its elements are divisible.

2.1.3 Analytic Spaces

Throughout this paper, k is a non-archimedean *analytic field*, i.e. a real-valued field which is complete with respect to its valuation. Trivial valuation is allowed. All analytic spaces we will consider are k -analytic spaces in the sense of [5, Sect. 1].

In addition, we fix a divisible subgroup $H \subseteq \mathbf{R}_{>0}$ such that $H \neq 1$ and $|k^\times| \subseteq H$ and consider only H -strict analytic spaces in the sense of [10]. For shortness, we will often call them *analytic spaces*.

2.1.4 The G -Topology

The usual topology of an analytic space X can be used for working with coherent sheaves only when X is good. In general, one has to work with the G -topology of analytic domains whose coverings are the set-theoretical coverings $U = \cup_{i \in I} U_i$ such that $\{U_i\}$ is a quasi-net on U in the sense of [4, Sect. 1.1]. These coverings are usually called G -coverings or *admissible* coverings. By X_G we denote the associated site: its objects are (H -strict) analytic domains and coverings are the G -admissible ones. The structure sheaf \mathcal{O}_{X_G} of X is a sheaf on X_G , and by $\mathcal{O}_{X_G}^\circ$ we denote the subsheaf of k° -algebras whose sections have spectral seminorm bounded by 1, i.e. $\mathcal{O}_{X_G}^\circ(V) = \mathcal{A}_V^\circ$ for an affinoid domain $V = \mathcal{M}(\mathcal{A}_V)$.

2.2 Seminormed Rings and Modules

In Sect. 2.2 we recall well-known facts and definitions concerning seminorms. All seminorms we consider are non-archimedean. Ring seminorms will be denoted $\|\cdot\|, \|\cdot\|'$, etc., and module seminorms will be denoted $\|\cdot\|, \|\cdot\|'$, etc.

2.2.1 Seminormed Abelian Groups

Throughout this paper, a *seminorm* on an abelian group A is a function $\|\cdot\| : A \rightarrow \mathbf{R}_{\geq 0}$ such that

- (1) $\|0\| = 0$,
- (2) the inequality $\|x - y\| \leq \max(\|x\|, \|y\|)$ holds for any $x, y \in A$.

Note that (2) is a short way to encode the more standard conditions that $\|x + y\| \leq \max(\|x\|, \|y\|)$ and $\|-x\| = \|x\|$. The pair $(A, \|\cdot\|)$ will be called a *seminormed group*. If $\|\cdot\|$ has trivial kernel, then it is called a *norm*.

2.2.2 Bounded and Non-expansive Homomorphisms

A homomorphism $\phi : A \rightarrow B$ is called *bounded* with respect to $\|\cdot\|_A$ and $\|\cdot\|_B$ if there exists $C = C(\phi)$ such that $\|a\| \leq C\|\phi(a)\|$ for any $a \in A$. If $C = 1$, then ϕ is called *non-expansive*.

2.2.3 The Non-expansive Category

As in [3, Sect. 5.1], the category of seminormed abelian groups with non-expansive homomorphisms will be called the *non-expansive category* (of seminormed abelian groups).

- Remark 2.2.4.* (i) Often one works with the larger category whose morphisms are arbitrary bounded morphisms; let us call it the *bounded category*. In fact, it is equivalent to the localization of the non-expansive category by bounded maps that possess a bounded inverse.
- (ii) Working with the bounded category is natural when one wants to study seminorms up to equivalence; for example, this is the case in the theory of Banach spaces. On the other side, working with the non-expansive category is natural when one distinguishes equivalent seminorms, so this fits the goals of the current paper.
- (iii) A serious advantage of working with the non-expansive category is that one can describe limits and colimits in a simple way, see [3, Sect. 5.1].

2.2.5 Seminormed Rings

A *seminorm* on a ring A is a seminorm $\|\cdot\|$ on the underlying group $(A, +)$ that satisfies $\|1\| = 1$ and $\|xy\| \leq \|x\|\|y\|$ for any x, y . A ring with a fixed seminorm is called a *seminormed ring*. Usually it will be denoted by calligraphic letters and the seminorm will be omitted from the notation, e.g. $\mathcal{A} = (\mathcal{A}, \|\cdot\|_{\mathcal{A}})$. An important example of a normed ring is a real-valued field.

2.2.6 Seminormed Modules

A seminormed \mathcal{A} -module M is an \mathcal{A} -module provided with a seminorm $\|\cdot\|_M$ such that $\|am\|_M \leq \|a\|_{\mathcal{A}}\|m\|_M$ for any $a \in \mathcal{A}$ and $m \in M$. The notions of non-expansive homomorphisms of seminormed rings and modules are defined in the obvious way.

2.2.7 Quotient Seminorms and Cokernels

If $\phi : A \rightarrow B$ is a surjective homomorphism of seminormed abelian groups, rings, or modules, then the *quotient seminorm* on B is defined by $\|x\|_B = \inf_{y \in \phi^{-1}(x)} \|y\|_A$. It is the maximal seminorm such that ϕ is non-expansive, hence, in the case of groups and modules, $(B, \|\cdot\|_B)$ is the cokernel of any non-expansive homomorphism $C \rightarrow A$ whose image is $\text{Ker}(\phi)$.

2.2.8 Strictly Admissible Homomorphisms

Recall that a homomorphism $\psi : C \rightarrow D$ of seminormed abelian groups (resp. rings or modules) is called *admissible* if the quotient seminorm on $\psi(C)$ is equivalent to the seminorm induced from D . In the non-expansive category, it is natural to consider the following more restrictive notion: ψ is *strictly admissible* if the quotient seminorm on $\psi(C)$ equals to the seminorm induced from D .

2.2.9 Tensor Products

Given a seminormed ring \mathcal{A} , by *tensor product* of seminormed \mathcal{A} -modules M and N we mean the module $L = M \otimes_{\mathcal{A}} N$ provided with the *tensor seminorm* $\| \cdot \|_{\otimes}$ such that $\|l\|_{\otimes} = \inf(\max_i \|m_i\| \cdot \|n_i\|)$, where the infimum is taken over all representations $l = \sum_{i=1}^n m_i \otimes n_i$. Note that $\| \cdot \|_{\otimes}$ is the maximal seminorm such that the bilinear map $\phi : M \times N \rightarrow L$ is *non-expansive*, i.e. satisfies $\|\phi(m, n)\|_{\otimes} \leq \|m\| \cdot \|n\|$. Obviously, $M \times N \rightarrow (L, \| \cdot \|_{\otimes})$ is the universal non-expansive bilinear map.

2.2.10 Exterior and Symmetric Powers

If M is a seminormed \mathcal{A} -module, then the modules $S^n M$ and $\bigwedge^n M$ acquire a natural seminorm as follows: both are quotients of $\otimes^n M$, so we consider the tensor seminorm on $\otimes^n M$ and endow $S^n M$ and $\bigwedge^n M$ with the quotient seminorms. The latter can be characterized as the maximal seminorms such that $\|m_1 \otimes \dots \otimes m_n\|_{S^n}$ and $\|m_1 \wedge \dots \wedge m_n\|_{\bigwedge^n}$ do not exceed $\prod_{i=1}^n \|m_i\|$ for any $m_1, \dots, m_n \in M$.

2.2.11 Filtered Colimits

Assume that $\{(\mathcal{A}_{\lambda}, \| \cdot \|_{\lambda})\}_{\lambda \in \Lambda}$ is a filtered family of seminormed abelian groups (resp. rings or \mathcal{A} -modules) with non-expansive transition homomorphisms, $\mathcal{A} = \text{colim}_{\lambda} \mathcal{A}_{\lambda}$ is the filtered colimit, and $f_{\lambda} : \mathcal{A}_{\lambda} \rightarrow \mathcal{A}$ are the natural maps. We endow \mathcal{A} with the *colimit seminorm* given by $\|a\| = \inf_{\lambda \in \Lambda, b \in f_{\lambda}^{-1}(a)} \|b\|_{\lambda}$. Obviously, this is the maximal seminorm making each homomorphism f_{λ} non-expansive, and hence \mathcal{A} is the colimit of \mathcal{A}_{λ} in the category of seminormed rings.

Lemma 2.2.12. *Filtered colimits of seminormed rings and modules are compatible with quotients, tensor products, symmetric and exterior powers.*

Proof. For concreteness, consider the case of tensor products. Assume that $\{M_i\}$ and $\{N_i\}$ are two filtered families and set $L_i = M_i \otimes N_i$, $M = \text{colim}_i M_i$, $N = \text{colim}_i N_i$ and $L = \text{colim}_i L_i$. By the universal properties of tensor and colimit seminorms we obtain a non-expansive homomorphism $\phi : L \rightarrow M \otimes N$, whose underlying homomorphism is the classical isomorphism. It remains to show that ϕ^{-1} is non-expansive too. If $l \in M \otimes N$ satisfies $\|l\|_{M \otimes N} < r$, then $l = \sum_{j=1}^r m_j \otimes n_j$ with $\max_j (\|m_j\|_M \cdot \|n_j\|_N) < r$. Choosing i large enough we can achieve that m_j and n_j come from elements $m'_j \in M_i$ and $n'_j \in N_i$ whose norms are so close to the norms of m_j and n_j that the inequality $\max_j (\|m'_j\|_{M_i} \cdot \|n'_j\|_{N_i}) < r$ holds. Then $l' = \sum_j m'_j \otimes n'_j \in L_i$ is a lifting of l satisfying $\|l'\|_{L_i} < r$ and hence $\|\phi^{-1}(l)\|_L < r$, as required. \square

2.2.13 The Completion Along a Seminorm

Throughout this paper “complete” means what one sometimes calls “Hausdorff and complete.” Similarly, by “completion” we mean what one sometimes calls “separated completion.”

The completion of A with respect to the semimetric $d(x, y) = \|x - y\|$ is denoted \widehat{A} . Note that $\|\cdot\|$ extends to \widehat{A} by continuity and the completion map $\alpha : A \rightarrow \widehat{A}$ is an *isometry* (i.e., $\|x\| = \|\alpha(x)\|$) whose kernel is the kernel of $\|\cdot\|$. Completion is functorial and, in the case of groups or modules, it takes exact sequences to *semiexact* sequences, i.e. sequences in which $\text{Im}(d_{i+1})$ is dense in $\text{Ker}(d_i)$. Moreover, if all morphisms of an exact sequence are strictly admissible, then the completion is exact and strictly admissible.

2.2.14 Banach Rings and Modules

A seminormed ring or module is called *Banach* if it is complete; in particular, the seminorm is a norm. Note that a Banach \mathcal{A} -module M is automatically a Banach $\widehat{\mathcal{A}}$ -module.

2.2.15 Unit Balls

The unit ball of a seminormed ring \mathcal{A} will be denoted \mathcal{A}^\diamond ; it is a subring of \mathcal{A} . The unit ball of a seminormed \mathcal{A} -module M will be denoted M^\diamond ; it is an \mathcal{A}^\diamond -module.

Remark 2.2.16. (i) Perhaps, M° would be a better notation, but the sign \circ is traditionally reserved for the unit ball of the spectral seminorm of a Banach algebra. So, we use the diamond sign \diamond instead.

(ii) If k is a real-valued field, $0 \neq \pi \in k^{\circ\circ}$, and V is a seminormed k -space, then the induced topology on V^\diamond is the π -adic one. In this case, V is Banach if and only V^\diamond is π -adic.

2.2.17 Bounded Categories

For the sake of completeness, we make some remarks on the categories of seminormed abelian groups (resp. rings or \mathcal{A} -modules) with bounded homomorphisms. Usually, seminormed \mathcal{A} -modules in our sense are called non-expansive and a general seminormed \mathcal{A} -module M is defined to be an \mathcal{A} -module provided with a seminorm $\|\cdot\|_M$ such that there exists $C = C(M)$ satisfying $\|am\|_M \leq C|a|_{\mathcal{A}}\|m\|_M$ for any $a \in \mathcal{A}$ and $m \in M$. Some results below can be extended to bounded homomorphisms using the following lemma, but we will not pursue this direction in the sequel.

Lemma 2.2.18. *Assume that $\phi : S \rightarrow R$ is a bounded homomorphism of seminormed abelian groups, rings, or modules. Then replacing the norm of S with an equivalent norm one can make ϕ non-expansive.*

Proof. Define a new seminorm $\| \cdot \|'_S$ by $\|a\|'_S = \max(\|a\|_S, \|\phi(a)\|_R)$. □

2.3 K° -Modules

In this section we study K° -modules for a real-valued field K . This material will be heavily used later, in particular, because such modules appear as unit balls of seminormed K -vector spaces.

2.3.1 Almost Isomorphisms

Let K be a real-valued field and let M be a K° -module. We say that M is a *torsion* module if any its element has a non-zero annihilator. By M_{tor} and $M_{\text{tf}} = M/M_{\text{tor}}$ we denote the (maximal) torsion submodule and the (maximal) torsion free quotient, respectively. We say that an element $x \in M$ is *almost zero* if for any $r < 1$ there exists $\pi \in K^\circ$ such that $r < |\pi|$ and $\pi x = 0$, and x is called *essential* otherwise. If $|K^\times|$ is discrete, then any non-zero element is essential. As in [18], we say that a module is almost zero if all its elements are so, and a homomorphism is an *almost isomorphism* if its kernel and cokernel almost vanish.

2.3.2 Almost Isomorphic Envelope

Let N be a K° -module with submodules M and M' . We say that M and M' are *almost isomorphic* as submodules if the embeddings $M \hookrightarrow M + M'$ and $M' \hookrightarrow M + M'$ are almost isomorphisms. By the *almost isomorphic envelope*, or just *envelope*, M^{env} of M in N we mean the maximal submodule $M' \subseteq N$ which is almost isomorphic to M . Obviously, M' consists of all elements $x \in N$ such that for any $r < 1$ there exists $\pi \in K^\circ$ satisfying $r < |\pi|$ and $\pi x \in M$. Thus, if $|K^\times|$ is discrete then $M^{\text{env}} = M$ for any M and N , and if $|K^\times|$ is dense then M^{env} is the maximal submodule $M' \subseteq N$ such that $K^{\circ\circ}M' \subseteq M$. For example, if $N = K$ and $|K^\times|$ is dense, then $(K^{\circ\circ})^{\text{env}} = K^\circ$.

2.3.3 Adic Seminorm

Given a K° -module M we define the *adic seminorm* as

$$\|x\|_{\text{adic}} = \inf_{a \in K^\circ, x \in aM} |a|.$$

Lemma 2.3.4. *Let M be a K° -module and $x \in M$ an element.*

- (i) *The adic seminorm is the maximal K° -seminorm on M bounded by 1.*
- (ii) *$\|x\|_{\text{adic}} = 0$ if and only if x is divisible (see Sect. 2.1.2).*

Proof. The claim is almost obvious and the only case that requires a little care is when the valuation is trivial. In this case, $\|\cdot\|_{\text{adic}}$ is trivial, i.e. $\|x\|_{\text{adic}} = 1$ whenever $x \neq 0$, and so $\|x\|_{\text{adic}} = 0$ if and only if $x = 0$, i.e. x is divisible. \square

Lemma 2.3.5. (i) *Any homomorphism $\phi : M \rightarrow N$ between K° -modules is non-expansive with respect to the adic seminorms.*

- (ii) *Assume that M and N are torsion free. Then ϕ is an isometry if and only if $\text{Ker}(\phi)$ is divisible and $\text{Coker}(\phi)$ contains no essential torsion elements.*

Proof. The first claim is obvious, so let us prove (ii). If $Q = \text{Ker}(\phi)$ is not divisible, then it contains an element x which is not infinitely divisible in Q and hence also not infinitely divisible in M . Then $\|x\|_{\text{adic}} \neq 0$, and hence ϕ is not an isometry. So, we can assume that Q is divisible and we should prove that in this case ϕ is an isometry if and only if $\text{Coker}(\phi)$ contains no essential torsion elements. Since Q is divisible, M/Q is torsion free and the map $M \rightarrow M/Q$ is an isometry. Thus, replacing M with M/Q we can assume that ϕ is injective. In this case the assertion is obvious. \square

2.3.6 Finitely Presented Modules

The following result is proved as its classical analogue over DVR by reducing a matrix to a diagonal one by elementary operations.

Lemma 2.3.7. *Assume that $L \subseteq M$ are free K° -modules of ranks l and m , respectively. Then there exists a basis e_1, \dots, e_m of M and elements $\pi_1, \dots, \pi_l \in K^\circ$ such that $\pi_1 e_1, \dots, \pi_l e_l$ is a basis of L .*

Corollary 2.3.8. *Any finitely presented K° -module M is of the form $\bigoplus_{i=1}^n M_i$ with each M_i non-zero cyclic, say $M_i = K^\circ/\pi_i K^\circ$, and $1 > |\pi_1| \geq \dots \geq |\pi_n| \geq 0$. In addition, the sequence $|\pi_1|, \dots, |\pi_n|$ is determined by M uniquely.*

2.4 Seminorms on K -Vector Spaces

Now, let us study seminorms on vector spaces over a real-valued field K .

2.4.1 Orthogonal Bases and Cartesian Spaces

We recall some results from [8, Chap. 2]. For simplicity we only consider the finite-dimensional case; generalizations to the case of infinite dimension can be found in loc.cit. So, assume that $(V, \|\cdot\|)$ is a finite-dimensional seminormed K -vector

space. A basis e_1, \dots, e_n of V is called r -orthogonal, where $r \in (0, 1]$, if for any $v = \sum_i a_i e_i$ the inequality $\|v\| \geq r \max_i(|a_i| \cdot \|e_i\|)$ holds. If $r = 1$, then the basis is called orthogonal, and if in addition $\|e_i\| = 1$ for $1 \leq i \leq n$ then the basis is called orthonormal. One says that the seminormed space V and the seminorm $\| \cdot \|$ are weakly cartesian (resp. cartesian, resp. strictly cartesian) if V possesses an r -orthogonal basis for some r (resp. an orthogonal basis, resp. an orthonormal basis).

Remark 2.4.2. (i) If V is weakly cartesian, then it possesses an r -orthogonal basis for any $r < 1$, see [8, Proposition 2.6.2/3].

(ii) If K is complete, then V is weakly cartesian if and only if its seminorm is a norm, see [8, Proposition 2.3.3/4]. Conversely, if K is not complete, say $x \in \widehat{K} \setminus K$ then it is easy to see that $K + Kx$ with the norm induced from \widehat{K} is a normed vector space of dimension two which is not weakly cartesian (e.g., its completion is the one-dimensional \widehat{K} -vector space \widehat{K}).

(iii) If K is spherically complete, then any normed vector space is cartesian by [8, Proposition 2.4.4/2]. Conversely, if K is not spherically complete one can easily construct two-dimensional normed vector spaces which are not cartesian.

2.4.3 Index of Norms

If U is one-dimensional with basis e , then sending a seminorm to its value on e provides a one-to-one correspondence between seminorms (resp. norms) on V and the half-line $\mathbf{R}_{\geq 0}$ (resp. $\mathbf{R}_{> 0}$). In particular, if $\| \cdot \|$ is a norm and $\| \cdot \|'$ is a seminorm on U , then the index $[[\| \cdot \|' : \| \cdot \|] = \|e\|'/\|e\|$ is a well-defined number independent of the choice of e .

We can extend this construction using the top exterior powers. Assume that $\dim(V) = d$ and set $U = \det(V) = \bigwedge^d V$. Any seminorm $\| \cdot \|$ on V induces the determinant seminorm $\| \cdot \|_{\det} = \| \cdot \|_{\bigwedge^d}$ on U , see Sect. 2.2.10. Moreover, if $\| \cdot \|$ is weakly cartesian then it is easy to see that $\| \cdot \|_{\det}$ is a norm, hence for any pair of weakly cartesian norms we can define the index $[[\| \cdot \|' : \| \cdot \|]$ to be $[[\| \cdot \|'_{\det} : \| \cdot \|_{\det}]$. Obviously, the index is transitive, i.e.

$$[[\| \cdot \|' : \| \cdot \|] \cdot [[\| \cdot \|'' : \| \cdot \|'] = [[\| \cdot \|'' : \| \cdot \|] .$$

Remark 2.4.4. The intuitive meaning of the index is that it measures the inverse ratio of volumes of the unit balls of the norms, at least when $|K^\times|$ is dense. This will be made precise in Lemma 2.6.4 below.

2.5 Unit Balls

Next, we study seminorms in terms of their unit balls. In this section, the trivially valued case will be uninteresting: we will often have to exclude it and in the remaining cases it will reduce to a triviality.

2.5.1 Semilattices

Assume that V is a K -vector space and $M \subseteq V$ is a K° -submodule. If $M \otimes_{K^\circ} K = V$, then we say that M is a *semilattice* of V . If a semilattice M is a free K° -module, then we say that M is a *lattice*. Note that M is a semilattice if and only if it contains a basis of V and hence contains a lattice.

2.5.2 Almost Unit Balls

For any K -seminorm $\|\cdot\|$ on a K -vector space V the unit ball V^\diamond is a K° -module. More generally, we say that a K° -submodule $M \subseteq V$ is an *almost unit ball* of $\|\cdot\|$ if it is almost isomorphic to V^\diamond in the sense of Sect. 2.3.2. In particular, V^\diamond itself is an almost unit ball too. If the valuation on K is non-trivial, then any almost unit ball M is a semilattice.

2.5.3 Seminorm Determined by a Semilattice

Any semilattice $M \subseteq V$ determines a K -seminorm on V as follows: M possesses the canonical adic seminorm and there is a unique way to extend it to a K -seminorm $\|\cdot\|_M$ on $V = KM$. By Lemma 2.3.4, $\|\cdot\|_M$ is the maximal seminorm whose unit ball contains M . The following lemma shows to which extent these correspondences between seminorms and semilattices are inverse one to another. The proof is simple so we omit it.

Lemma 2.5.4. *Let V be a vector space over a real-valued field K .*

- (i) *Assume that $M \subseteq V$ is a semilattice and let $\|\cdot\|$ be the maximal seminorm such that $\|M\| \leq 1$. Then M^{env} is the unit ball of $\|\cdot\|$. In particular, M is an almost unit ball of V .*
- (ii) *Assume that the valuation on K is non-trivial, V is provided with a K -seminorm, $\|\cdot\|$ and V^\diamond is the unit ball. Let $\|\cdot\|'$ be the maximal seminorm bounded by 1 on V^\diamond . Then $\|\cdot\|'$ is the minimal seminorm that dominates $\|\cdot\|$ and takes values in the closure of $|K|$ in $\mathbf{R}_{\geq 0}$.*

Corollary 2.5.5. *Assume that K is a real-valued field whose valuation is non-trivial and V is a K -vector space. Then the above constructions establish a one-to-one correspondence between almost isomorphism classes of semilattices $M \subseteq V$ and K -seminorms on V taking values in the closure of $|K|$ in $\mathbf{R}_{\geq 0}$.*

2.5.6 Bounded Semilattices

Assume now that V is finite-dimensional and let us interpret the results of Sect. 2.4 in terms of semilattices. A semilattice M of V is called *bounded* if it is contained in a lattice. Equivalently, for some (and then any) sublattice $L \subseteq M$ one has that $\pi M \subseteq L$ for some $0 \neq \pi \in K^\circ$.

Lemma 2.5.7. *Let M be a semilattice in a finite-dimensional K -vector space V with the associated seminorm $\| \cdot \|_M$.*

- (i) *The following conditions are equivalent: (a) M is almost isomorphic to a lattice of V in the sense of Sect. 2.3.2, (b) M^{env} is a lattice, (c) $\| \cdot \|_M$ is strictly cartesian.*
- (ii) *The following conditions are equivalent: (d) M is bounded, (e) for any $\pi \in K^\circ$ there exists a lattice $L \subseteq M$ such that $\pi M \subseteq L$, (f) $\| \cdot \|_M$ is weakly cartesian.*

Proof. Since any lattice is an almost isomorphic envelope of itself, the equivalence (a) \iff (b) is clear. By Lemma 2.5.4(i), M^{env} is the unit ball of $\| \cdot \|_M$, hence a basis of V is orthonormal if and only if it is also a basis of M^{env} . This shows that (b) \iff (c).

Obviously, (e) implies (d). If $\pi M \subseteq L \subseteq M$ and $|\pi| = r$, then any basis of L is r -orthogonal with respect to $\| \cdot \|_M$ and we obtain that (d) \implies (f). Finally, let us prove that (f) implies (e). Assume that $\| \cdot \|_M$ is weakly cartesian and choose an r -orthogonal basis e_1, \dots, e_n for some $r \in (0, 1]$. If $0 < |\pi| < r$, then M is contained in the lattice $\bigoplus_{i=1}^n \pi^{-1} K^\circ e_i$. It follows that if $|K^\times|$ is discrete then M is a lattice, so we can assume that $|K^\times|$ is dense. Choose any number $s \in (0, 1)$. By Remark 2.4.2(i), there exists an s -orthogonal basis e_1, \dots, e_n . Multiplying e_i by elements of K we can also achieve that $e_i \in M$ and $e_i \notin \pi M$ for any $\pi \in K$ with $|\pi| < s$. Then e_i generate a lattice $L \subseteq M$ such that $\pi M \subseteq L$ for any $\pi \in K$ with $|\pi| < s^2$. \square

Remark 2.5.8. (i) By definition, $m \in M$ is divisible if and only if $\|m\|_M = 0$. So, Remark 2.4.2(ii) implies that K is complete if and only if any semilattice without non-zero divisible elements is bounded.

(ii) One can introduce a class of almost lattices using property (e) of Lemma 2.5.7 as the definition. We do not use this terminology since an almost lattice is the same as a bounded semilattice. Using Lemma 2.5.7 and Remark 2.4.2(iii) one can easily check that if K is spherically complete and $|K^\times| = \mathbf{R}_{>0}$ then M is an almost lattice if and only if it is almost isomorphic to a lattice, but in general there exist almost lattices which are not almost isomorphic to a lattice.

2.5.9 Index of Bounded Semilattices

If M is a semilattice in V , then $\bigwedge^i M$ is a semilattice in $\bigwedge^i V$. Furthermore, if M is bounded, then it is contained in a lattice L and hence $\bigwedge^i M$ is contained in the lattice $\bigwedge^i L$. In particular, $\bigwedge^i M$ is bounded. If $d = \dim(V)$, then we call $\det(M) = \bigwedge^d M$ the *determinant* of M . As in the case of seminorms, we define the *index* of two bounded semilattices M, N of V to be the ratio of their determinants:

$$[M : N] = |\det(M) : \det(N)| = \sup \{ |\pi| : \pi \in K, \pi \det(N) \subseteq \det(M) \}.$$

Lemma 2.5.10. *Let V be a finite-dimensional vector space over a real-valued field K , let M, M' be two bounded semilattices of V , and let $\|\cdot\|$ and $\|\cdot\|'$ be the associated norms on V . Then,*

- (i) *For any i , $\|\cdot\|_{\wedge^i}$ is the norm associated with $\wedge^i M$.*
- (ii) $[\|\cdot\| : \|\cdot\|'] = [M : M']^{-1}$.

Proof. Since M is an almost unit ball of $\|\cdot\|$, it is easy to see that $\|\cdot\|_{\wedge^i}$ is the maximal norm such that $\|v_1 \wedge \cdots \wedge v_i\|_{\wedge^i} \leq 1$ for any $v_1, \dots, v_i \in M$. Thus the module $\wedge^i M$, which is generated by the elements $v_1 \wedge \cdots \wedge v_i$ with $v_1, \dots, v_i \in M$, is an almost unit ball of $\|\cdot\|_{\wedge^i}$. This implies (i), and taking $i = \dim(V)$ we reduce (ii) to the one-dimensional case, which is clear. □

2.6 Content of K° -Modules

2.6.1 The Definition

Given a finitely presented K° -module M we represent it as in Corollary 2.3.8 and define the *content* of M to be $\text{cont}(M) = \prod_{i=1}^n |\pi_i|$. In general, we set $\text{cont}(M) = \inf_\alpha \text{cont}(M_\alpha)$, where M_α run through all finitely presented subquotients of M . Obviously, this is compatible with the definition in the finitely presented case. Note that $\text{cont}(M) = 0$ if M is not a torsion module.

Remark 2.6.2. The content invariant adequately measures the “size” of a torsion module M . In particular, $\text{cont}(M) = 1$ if and only if M almost vanishes. Also, if the valuation of K is discrete and π_K is a uniformizer then $\text{cont}(M) = |\pi_K|^{\text{length}(M)}$.

2.6.3 Relation to the Index

We will study the content by relating it to index of semilattices.

Lemma 2.6.4. *Assume that V is a finite-dimensional vector space over a real-valued field K and $L \subseteq M$ are two bounded semilattices of V . Then $[M : L] = \text{cont}(M/L)^{-1}$.*

Proof. If M and L are lattices, then the assertion follows by use of Lemma 2.3.7. In particular, this covers the case of DVRs, so in the sequel we assume that $|K^\times|$ is dense. Choose $0 \neq \pi \in K^\circ$. Then it is easy to see that $[M : \pi L] = \pi^{-d} [M : L]$ and $\text{cont}(M/\pi L) = \pi^d \text{cont}(M/L)$, where $d = \dim(V)$. Therefore, it suffices to prove the lemma for M and πL , and we can assume in the sequel that $L \subseteq \pi M$ for some π with $|\pi| < 1$.

Let π_1, π_2, \dots be elements of K° such that the sequence $|\pi_i|$ strictly increases and tends to 1. By Lemma 2.5.7(d) \iff (e), for each $i \in \mathbb{N}$ there exist lattices L_i and M_i such that

$$\pi_i L_i \subseteq L \subseteq L_i \subseteq M_i \subseteq M \subseteq \pi_i^{-1} M_i.$$

Since the case of lattices was established, it suffices to check the equalities

$$\lim_i \text{cont}(M_i/L_i) = \text{cont}(M/L), \quad \lim_i [M_i : L_i] = [M : L].$$

The latter follow from the observation that

$$\text{cont}(M_i/L_i) \geq \text{cont}(M/L) \geq \text{cont}(\pi_i^{-1} M_i / \pi_i L_i) = \pi_i^{2d} \text{cont}(M_i/L_i)$$

and

$$[M_i : L_i] \leq [M : L] \leq [\pi_i^{-1} M_i : \pi_i L_i] = \pi_i^{-2d} [M_i : L_i].$$

□

2.6.5 Properties

The following continuity result reduces computation of contents to the finitely generated case.

Lemma 2.6.6. *If a K° -module M is a filtered union of submodules M_i , then*

$$\text{cont}(M) = \inf_i \text{cont}(M_i) = \lim_i \text{cont}(M_i).$$

Proof. Any subquotient of M_i is a subquotient of M hence $\text{cont}(M_i) \geq \text{cont}(M)$. Also, it is easy to see that any finitely presented subquotient of M is a subquotient of each M_i with a large enough i . Hence the equalities hold. □

Now we can establish the main property of content, the multiplicativity.

Theorem 2.6.7. *If $0 \rightarrow M' \rightarrow M \rightarrow M'' \rightarrow 0$ is a short exact sequence of K° -modules, then $\text{cont}(M) = \text{cont}(M') \cdot \text{cont}(M'')$.*

Proof. If M is not torsion, then either M' or M'' is not torsion and hence both $\text{cont}(M)$ and $\text{cont}(M')\text{cont}(M'')$ vanish. So assume that M is torsion. First, we consider the case when M is finitely generated. Fix an epimorphism $L = (K^\circ)^n \twoheadrightarrow M$ and denote its kernel Q . Then L is a lattice in $V = K^n$ and since $M = L/Q$ is torsion, Q is a semilattice of V . Let Q' be the kernel of the composition $L \rightarrow M \rightarrow M''$ then $M'' = L/Q'$ and $M' = Q'/Q$. Hence the claim follows from the multiplicativity of the index and Lemma 2.6.4.

Assume now that M is a general torsion module and let $\{M_i\}$ be the family of finitely generated submodules of M . Then M' is the filtered union of the submodules $M'_i = M' \cap M_i$ and M'' is the filtered union of the modules $M''_i = M_i/M'_i$. By the above case, $\text{cont}(M_i) = \text{cont}(M'_i) \cdot \text{cont}(M''_i)$ and it remains to pass to the limit and use Lemma 2.6.6. \square

3 Metrization of Sheaves

3.1 Seminormed Sheaves

Throughout Sect. 3.1, \mathcal{C} is a site, i.e. a category provided with a Grothendieck topology. In our applications, \mathcal{C} will be the category associated with a G -topological space, more concretely, it will be of the form X_G for an analytic space X .

3.1.1 Quasi-Norms

The sup seminorm can be infinite on a non-compact set, so it is technically convenient to introduce the following notion. Let $\mathbf{R}_{\geq 0} = \mathbf{R}_{\geq 0} \cup \{\infty\}$ be the one-pointed compactification of $\mathbf{R}_{\geq 0}$ with addition and multiplication satisfying all natural rules and the rule $0 \cdot \infty = 0$. *Quasi-norms* on abelian groups, rings, and modules are defined similarly to seminorms but with the target $\mathbf{R}_{\geq 0}$. For example, a quasi-norm $\|\cdot\|$ on a ring A is a map $\|\cdot\| : A \rightarrow \mathbf{R}_{\geq 0}$ such that $\|0\| = 0$, $\|1\| = 1$, $\|a - b\| \leq \max(\|a\|, \|b\|)$ and $\|ab\| \leq \|a\| \cdot \|b\|$. The material of Sect. 2.2, including the constructions of quotient and tensor product seminorms, extends to quasi-norms straightforwardly.

3.1.2 Seminorms on Sheaves

Let \mathcal{A} be a sheaf of abelian groups on \mathcal{C} . By a *pre-quasi-norm* $\|\cdot\|$ on \mathcal{A} we mean a family of quasi-norms $\|\cdot\|_U$ on $\mathcal{A}(U)$, where U runs through the objects of \mathcal{C} , such that the restriction maps $\mathcal{A}(U) \rightarrow \mathcal{A}(V)$ are non-expansive. In the same way one defines quasi-norms on sheaves of rings and quasi-norms on sheaves of modules over a sheaf of rings provided with a pre-quasi-norm. The constructions we describe below for pre-quasi-norms on sheaves of abelian groups hold also for pre-quasi-norms on sheaves of rings and modules.

A pre-quasi-norm is called a *quasi-norm* if it satisfies the following locality condition: for any covering $\{U_i \rightarrow U\}$ in \mathcal{C} the equality $\|s\|_U = \sup_i \|s_i\|_{U_i}$ holds. A pre-quasi-norm is called *locally bounded* if for any U in \mathcal{C} and $s \in \mathcal{A}(U)$ there exists a covering $\{U_i \rightarrow U\}$ such that $\|s_i\|_{U_i} < \infty$ for any i , where s_i denotes $s|_{U_i}$.

A locally bounded quasi-norm will be called a *seminorm*, and once a seminorm $\| \cdot \|$ is fixed we call the pair $\mathcal{A} = (\mathcal{A}, \| \cdot \|)$ a *seminormed sheaf of abelian groups*.

Remark 3.1.3. (i) Our ad hoc definitions have the following categorical interpretation. If $\| \cdot \|$ is a pre-quasi-norm on \mathcal{A} , then the pair $(\mathcal{A}, \| \cdot \|)$ can also be viewed as a presheaf of quasi-normed abelian groups on \mathcal{C} , and this presheaf is a sheaf if and only if $\| \cdot \|$ is a quasi-norm.

(ii) If an object U is quasi-compact (i.e., any of its coverings possesses a finite refinement) and $\| \cdot \|$ is locally bounded, then $\| \cdot \|_U$ is a seminorm. In particular, if the subcategory of quasi-compact objects \mathcal{C}_c is cofinal in \mathcal{C} , then seminorms on \mathcal{A} can be viewed as sheaves of seminormed abelian groups on \mathcal{C}_c with the underlying sheaf of abelian groups \mathcal{A} .

3.1.4 Sheafification

For any pre-quasi-norm $\| \cdot \|'$ on a sheaf of abelian groups \mathcal{A} we define the *sheafification* $\| \cdot \| = \alpha(\| \cdot \|')$ by the rule

$$\|s\|_U = \inf_{\{U_i \rightarrow U\}} \sup_{i \in I} \|s_i\|'_{U_i},$$

where U is an object of \mathcal{C} , $s \in \mathcal{A}(U)$ and the infimum is over all coverings of U .

Lemma 3.1.5. *Keep the above notation. Then $\| \cdot \|$ is a quasi-norm and it is the maximal quasi-norm dominated by $\| \cdot \|'$. In addition, if $\| \cdot \|'$ is locally bounded, then $\| \cdot \|$ is a seminorm.*

Proof. The fact that $\| \cdot \|$ is a quasi-norm follows easily from the transitivity of coverings. The other assertions are obvious. □

Remark 3.1.6. The universal property from the first part of the lemma justifies the notion “sheafification.” In fact, the same argument shows that $(\mathcal{A}, \| \cdot \|') \rightarrow (\mathcal{A}, \| \cdot \|)$ is the universal (non-expansive) map from $(\mathcal{A}, \| \cdot \|')$ to a sheaf of abelian groups with a quasi-norm. Thus, $(\mathcal{A}, \| \cdot \|)$ is even the sheafification of $(\mathcal{A}, \| \cdot \|')$ as a sheaf of quasi-normed abelian groups.

3.1.7 Operations on Seminormed Sheaves

Let us extend various operations, including quotients, tensor products, symmetric powers, and exterior powers to seminormed sheaves. This is done in two stages. First one works with sections over each U separately. This produces a sheaf with a pre-quasi-norm, which is easily seen to be locally bounded. Then one sheafifies this pre-quasi-norm if needed. As in the case of usual sheaves of rings and modules, the second step is needed in the case of constructions that are not compatible with limits, such as colimits (including quotients), tensor products, etc.

For the sake of illustration, let us work this out in the case of quotients. Assume that $\phi : \mathcal{A} \rightarrow \mathcal{B}$ is an epimorphism of sheaves of abelian groups and \mathcal{A} is provided with a seminorm $\| \cdot \|$. First, we endow \mathcal{B} with the quotient pre-quasi-norm $\| \cdot \|'$, i.e. for each U in \mathcal{C} the quasi-norm $\| \cdot \|'_U$ on $\mathcal{B}(U)$ is the maximal quasi-norm making the map $(\mathcal{A}(U), \| \cdot \|_U) \rightarrow (\mathcal{B}(U), \| \cdot \|'_U)$ non-expansive. Then the quotient seminorm $\| \cdot \|'$ on \mathcal{B} is defined to be the sheafification of $\| \cdot \|'$. In particular, $\| \cdot \|'$ is the maximal seminorm on \mathcal{B} making the homomorphism $(\mathcal{A}, \| \cdot \|) \xrightarrow{\phi} (\mathcal{B}, \| \cdot \|')$ non-expansive.

Remark 3.1.8. Provide $\mathcal{K} = \text{Ker}(\phi)$ with any seminorm $\| \cdot \|_{\mathcal{K}}$ making the embedding $\mathcal{K} \hookrightarrow \mathcal{A}$ non-expansive. For example, one can take the restriction of $\| \cdot \|$ on \mathcal{K} . Then the universal property characterizing ϕ implies that $(\mathcal{B}, \| \cdot \|')$ is the cokernel of the map of the seminormed sheaves of abelian groups $(\mathcal{K}, \| \cdot \|_{\mathcal{K}}) \rightarrow (\mathcal{A}, \| \cdot \|)$.

3.1.9 Pushforwards

Recall that a morphism $f : \mathcal{C}' \rightarrow \mathcal{C}$ of sites is a functor $\mathcal{F} : \mathcal{C} \rightarrow \mathcal{C}'$ that satisfies certain properties, see [36, Tag:00X0]. If \mathcal{A}' is a sheaf of abelian groups on \mathcal{C}' with a pre-quasi-norm $\| \cdot \|'$, then we endow $\mathcal{A} = f_*(\mathcal{A}')$ with the *pushforward pre-quasi-norm* $\| \cdot \|$ as follows: for any U in \mathcal{C} and $U' = \mathcal{F}(U)$ we have that $\mathcal{A}(U) = \mathcal{A}'(U')$ and we set $\| \cdot \|_U = \| \cdot \|_{U'}$. Clearly, if $\| \cdot \|'$ is a quasi-norm, then $\| \cdot \|$ is also a quasi-norm, but the pushforward of a seminorm can be unbounded.

Remark 3.1.10. An important case when the pushforward of a seminorm is a seminorm is when f corresponds to a proper map of topological spaces. This fact will not be used, so we do not check it here.

3.1.11 Pullbacks

Assume, now, that \mathcal{A} is a sheaf of abelian groups on \mathcal{C} with a pre-quasi-norm $\| \cdot \|$ and $\mathcal{A}' = f^{-1}(\mathcal{A})$. Then $\mathcal{A}'(U') = \text{colim}_{U' \rightarrow \mathcal{F}(U)} \mathcal{A}(U)$ and we define $\| \cdot \|'_{U'}$ to be the colimit of the quasi-norms $\| \cdot \|_U$. If $\| \cdot \|$ is a quasi-norm, it still may happen that the pre-quasi-norm $\| \cdot \|'$ is not a quasi-norm, so we define $f^{-1}(\| \cdot \|)$ to be the sheafification of $\| \cdot \|'$. On the positive side we note that local boundedness is preserved by the pullback, so if $\| \cdot \|$ is a seminorm then $f^{-1}(\| \cdot \|)$ is a seminorm too.

3.2 Points and Stalks of Seminorms

In this section we study stalks of seminorms at points and describe seminorms that are fully controlled by stalks. For concreteness, we usually consider sheaves of abelian groups but the cases of rings and modules are similar.

3.2.1 Points of Sites

We use the terminology of [36, Tag:00Y3] when working with points of sites. Recall that a point x of \mathcal{C} is a functor $x : \mathcal{C} \rightarrow \mathbf{Sets}$ that satisfies certain properties and given an object U of \mathcal{C} an elements $f \in x(U)$ is interpreted as a morphism $f : x \rightarrow U$. The set of isomorphism classes of points of \mathcal{C} will be denoted $|\mathcal{C}|$.

3.2.2 Stalks

Assume that $\| \cdot \|$ is a pre-quasi-norm on a sheaf of abelian groups \mathcal{A} . Given a point $x \in |\mathcal{C}|$, we endow the stalk $\mathcal{A}_x = \text{colim}_{x \rightarrow U} \mathcal{A}(U)$ at x with the *stalk quasi-norm* $\| \cdot \|_x$ defined as follows: if $x \rightarrow V$ and $s \in \mathcal{A}(V)$, then $\|s\|_x = \inf_{x \rightarrow U \rightarrow V} \|s\|_U$. The following result follows by unveiling the definitions.

Lemma 3.2.3. *Assume that $\| \cdot \|$ is a pre-quasi-norm on a sheaf of abelian groups \mathcal{A} . Then,*

- (i) *The stalks of $\| \cdot \|$ and of the sheafification $\alpha(\| \cdot \|)$ coincide.*
- (ii) *If $\| \cdot \|$ is locally bounded, then the stalks $\| \cdot \|_x$ are seminorms.*

Also, Lemma 2.2.12 implies the following result.

Lemma 3.2.4. *Stalks of seminormed sheaves of rings or modules are compatible with quotients, tensor products, symmetric and exterior powers. For example, given a seminormed sheaf of rings \mathcal{A} , seminormed \mathcal{A} -modules \mathcal{M}, \mathcal{N} and $\mathcal{L} = \mathcal{M} \otimes_{\mathcal{A}} \mathcal{N}$, and a point $x \in |\mathcal{C}|$, there is a natural isomorphism of seminormed \mathcal{A}_x -modules $\mathcal{M}_x \otimes_{\mathcal{A}_x} \mathcal{N}_x = \mathcal{L}_x$.*

3.2.5 Semicontinuity

Let $\mathcal{P} \subseteq |\mathcal{C}|$ be a subset. A seminorm $\| \cdot \|$ on \mathcal{A} induces the set of stalk seminorms $\{\| \cdot \|_x\}_{x \in \mathcal{P}}$ which satisfies the following semicontinuity condition: if $x \in \mathcal{P}$ and $s_x \in \mathcal{A}_x$, then for any $\varepsilon > 0$ there exists a morphism $x \rightarrow U$ such that s_x is induced from $s \in \mathcal{A}(U)$ and $\|s\|_U < \|s\|_x + \varepsilon$, in particular, $\|s\|_y < \|s\|_x + \varepsilon$ for any $y \rightarrow U$. Any family of seminorms satisfying this condition will be called *upper semicontinuous*.

Example 3.2.6. If \mathcal{C} is the site of open subsets of a topological space X , then a family of seminorms $\{\| \cdot \|_x\}_{x \in X}$ is upper semicontinuous if and only if for any open U and a section $s \in \mathcal{A}(U)$ the function $\|s\| : U \rightarrow \mathbf{R}$ sending x to $\|s\|_x$ is upper semicontinuous.

3.2.7 \mathcal{P} -Seminorms

Let $\mathcal{P} \subseteq |\mathcal{C}|$. Given a seminorm $\| \cdot \|$ on \mathcal{A} set $\|s\|_{\mathcal{P},U} = \sup_{x \in \mathcal{P}, x \rightarrow U} \|s_x\|_x$. Clearly, $\| \cdot \|_{\mathcal{P}}$ is a seminorm on \mathcal{A} , and we say that $\| \cdot \|$ is a \mathcal{P} -seminorm if $\| \cdot \| = \| \cdot \|_{\mathcal{P}}$.

Lemma 3.2.8. *Keep the above notation then,*

- (i) $\| \cdot \|_{\mathcal{P}}$ is the maximal \mathcal{P} -seminorm on \mathcal{A} that is dominated by $\| \cdot \|$. The stalks of $\| \cdot \|_{\mathcal{P}}$ and $\| \cdot \|$ at any point of \mathcal{P} coincide.
- (ii) There is a natural bijection between \mathcal{P} -seminorms on \mathcal{A} and upper semicontinuous families of seminorms on the stalks of \mathcal{A} at the points of \mathcal{P} .

Proof. The first claim is clear. We noticed earlier that any seminorm gives rise to an upper semicontinuous family. Conversely, to any family of seminorms $\{\| \cdot \|_x\}$ on the stalks at the points $x \in \mathcal{P}$ we assign the sup seminorm $\|s\|_{\mathcal{P},U} = \sup_{x \in \mathcal{P}, x \rightarrow U} \|s_x\|_x$. To prove (ii) we should prove that if the family is upper semicontinuous then the stalk $\| \cdot \|_{\mathcal{P},x}$ coincides with the original seminorm $\| \cdot \|_x$. Clearly, $\| \cdot \|_x \leq \| \cdot \|_{\mathcal{P},x}$. In addition, if $\|s\|_x < r$, then there exists a morphism $x \rightarrow U$ such that s is induced from an element $s_U \in \mathcal{A}(U)$ and $\|s_U\|_y < r$ for any $y \rightarrow U$. Then $\|s\|_{\mathcal{P},x} \leq \|s_U\|_{\mathcal{P},U} < r$, thus proving (ii). \square

Remark 3.2.9. The class of \mathcal{P} -seminorms is not preserved under quotients and other operations. One can easily construct such examples already when \mathcal{C} is associated with the topological space $X = \{\eta, s\}$ consisting of an open point η and a closed point s , and $\mathcal{P} = \{\eta\}$.

3.2.10 Conservative Families

As one might expect, if \mathcal{C} possesses enough points then the theory of seminorms can be developed in terms of stalks. More concretely, assume that \mathcal{C} possesses a conservative family of points \mathcal{P} . Then seminorms and operations on them are completely controlled by the \mathcal{P} -stalks. In particular, one can define all operations on seminormed sheaves stalkwise. This follows from the following result.

Theorem 3.2.11. *Assume that \mathcal{P} is a conservative family of points of a site \mathcal{C} and \mathcal{A} is a sheaf of abelian groups (resp. rings, resp. modules over a seminormed sheaf of rings) on \mathcal{C} . Then any seminorm $\| \cdot \|$ on \mathcal{A} is a \mathcal{P} -seminorm. In particular, there is a natural bijective correspondence between seminorms on \mathcal{A} and upper semicontinuous families of seminorms on the stalks of \mathcal{A} at the points of \mathcal{P} .*

Proof. All cases are proved similarly, so assume that \mathcal{A} is a sheaf of abelian groups. Fix $r > 0$ and let $\mathcal{A}_r^\diamond \subseteq \mathcal{A}$ denote the ball of radius r ; it is the presheaf such that $\mathcal{A}_r^\diamond(U)$ is the set of all elements $s \in \mathcal{A}(U)$ with $\|s\|_U \leq r$. The locality condition satisfied by $\| \cdot \|$ implies that \mathcal{A}_r^\diamond is, in fact, a sheaf. Note that the stalk of \mathcal{A}_r^\diamond at a point x coincides with the ball $(\mathcal{A}_x)_r^\diamond$.

Define a subsheaf $\mathcal{A}_{\mathcal{P},r}^\diamond \xrightarrow{i_r} \mathcal{A}_r^\diamond$ by the condition that $s \in \mathcal{A}_{\mathcal{P},r}^\diamond(U)$ if for any $x \rightarrow U$ with $x \in \mathcal{P}$ the inequality $\|s\|_x \leq r$ holds. Note that $\mathcal{A}_{\mathcal{P},r}^\diamond$ is the r -ball of the seminorm $\|\cdot\|_{\mathcal{P}}$ and the stalk of $\mathcal{A}_{\mathcal{P},r}^\diamond$ at a point $x \in \mathcal{P}$ equals to $(\mathcal{A}_x)_r^\diamond$. So, the embedding of sheaves i_r induces isomorphisms of stalks at the points of \mathcal{P} , and using that \mathcal{P} is conservative we obtain that i_r is an equality. Thus, the balls of the seminorms $\|\cdot\|$ and $\|\cdot\|_{\mathcal{P}}$ coincide and hence the seminorms coincide. We proved that $\|\cdot\|$ is a \mathcal{P} -seminorm, and the second claim follows from Lemma 3.2.8(ii). \square

3.3 Sheaves on Analytic Spaces

In this section we study seminorms on \mathcal{O}_{X_G} -modules, where X is a H -strict k -analytic space and X_G denotes the site of H -strict analytic domains in X .

3.3.1 Analytic Points

Note that any point $x \in X$ defines a point of X_G , so we can view $|X|$ as a subset of $|X_G|$. A point of X_G is called *analytic* if its isomorphism class lies in $|X|$. A seminorm on a sheaf of abelian group on X_G is called G -analytic if it is an $|X|$ -seminorm in the sense of Sect. 3.2.7. For example, the spectral seminorm $|\cdot|$ on \mathcal{O}_{X_G} satisfies $|f|_U = \sup_{x \in U} |f(x)|$, i.e. it is G -analytic.

Example 3.3.2. (i) A typical example of a non-analytic point z on a unit disc $E = \mathcal{M}(k\{t\})$ is as follows: the family of neighborhoods of z is the set of all domains $U \subseteq E$ that contain an open annulus $r < |t| < 1$ for some r . One can view z as the maximal point with $|t|_z < 1$. In Huber adic geometry (see [21]) it corresponds to a valuation of height two such that $|t|_z < 1$ and $r < |t|_z$ for any real $r < 1$.

(ii) Using z one easily constructs a seminorm $\|\cdot\|$ on \mathcal{O}_{X_G} which is not G -analytic. For example, set $\|f\|_U = 0$ if $z \notin U$ and $\|f\|_U = \|f(p)\|$ otherwise, where p is the maximal point of E . In fact, all stalks of $\|\cdot\|$, excluding the stalk at z , are zero.

For the sake of simplicity, we only use analytic points in Sect. 3.3 (and in the main part of this paper). Arbitrary points of X_G will be described in Sect. 9.1, and then we will extend to them some results of this section.

3.3.3 Analytic Seminorms

Let \mathcal{A} be a sheaf of abelian groups on X_G with a seminorm $\|\cdot\|$. For any domain U and a section $s \in \mathcal{A}(U)$ consider the function $\|s\| : U \rightarrow \mathbf{R}_{\geq 0}$ that sends $x \in U$ to $\|s\|_x$. As we saw in Sect. 3.2.5, this function is upper G -semicontinuous, i.e. if $x \in X$

satisfies $\|s\|_x < r$ then there exists an analytic domain V such that $x \in V \subseteq U$ and $\|s\|_y < r$ for any $y \in V$. We say that the seminorm $\| \cdot \|$ is *analytic* if it is G -analytic and all functions $\|s\|$ are upper semicontinuous with respect to the usual topology of X , i.e. in the above situation V can be chosen to be a neighborhood of x .

- Remark 3.3.4.* (i) Both conditions are essential in the definition of analytic seminorms. For example, the seminorm $\| \cdot \|$ in Example 3.3.2 has zero stalks at all analytic points, hence all functions $\|s\|$ vanish though the seminorm is not analytic.
- (ii) Assume that $X = \cup_i X_i$ is an admissible covering. Then a function $X \rightarrow \mathbf{R}$ is upper semicontinuous if and only if its restrictions to X_i are upper semicontinuous. Therefore, a seminorm on \mathcal{A} is analytic if and only if its restrictions onto $\mathcal{A}|_{X_i}$ are analytic, i.e. analyticity is a G -local condition.
- (iii) The following simple observation will not be used, so we omit a justification. A seminorm $\| \cdot \|$ is analytic if and only if for any affinoid domain $U \subseteq X$, a section $s \in \mathcal{A}(U)$, a point $x \in U$ and $r > \|s\|_x$ there exists a neighborhood V of x in U such that $\|s\|_V < r$. This can also be reformulated using stalks of \mathcal{A} in the usual topology: $\| \cdot \|$ is analytic if and only if for any affinoid domain $U \subseteq X$ and a point $x \in U$ the map $\mathcal{A}|_{U|,x} \rightarrow \mathcal{A}_x$ is an isometry with respect to the stalk seminorms of $\| \cdot \|$, where $\mathcal{A}|_{U|}$ denotes the restriction of \mathcal{A} onto the topological space $|U|$.

3.3.5 Local Rings of \mathcal{O}_{X_G}

In Berkovich geometry, one usually works with local rings $\mathcal{O}_{X,x}$ of good spaces and their residue fields $\kappa(x) = \mathcal{O}_{X,x}/m_x$. We will also need the local rings $\mathcal{O}_{X_G,x}$ and their residue fields $\kappa_G(x) = \mathcal{O}_{X_G,x}/m_{G,x}$.

Lemma 3.3.6. *Let X be an analytic space with a point $x \in X$. Then*

- (i) $(\mathcal{O}_{X_G,x}, | \cdot |_x)$ is a seminormed local ring whose maximal ideal $m_{G,x}$ is the kernel of $| \cdot |_x$.
- (ii) The residue norm on $\kappa_G(x)$ is a valuation and the completion coincides with $\mathcal{H}(x)$. In particular, $\mathcal{H}(x)$ is the completion of $(\mathcal{O}_{X_G,x}, | \cdot |_x)$.

Proof. Note that $\mathcal{O}_{X_G,x}$ is the filtered colimit of the rings $\mathcal{O}_{V,x}$, where V is an affinoid domain containing x , and the transition maps are local. Hence (i) follows from the observation that the maximal ideal of $\mathcal{O}_{V,x}$ is the kernel of the restriction of $| \cdot |_x$. Furthermore, this implies that the residue field $\kappa_G(x)$ of $\mathcal{O}_{X_G,x}$ is the filtered union of the residue fields of $\mathcal{O}_{V,x}$. The latter are dense subfields of $\mathcal{H}(x)$, hence the same is true for $\kappa_G(x)$. It remains to note that the completion of $(\mathcal{O}_{X_G,x}, | \cdot |_x)$ is isomorphic to the completion of its quotient by the kernel of the seminorm. \square

3.3.7 Seminormed \mathcal{O}_{X_G} -Modules

Assume that $(\mathcal{F}, \|\cdot\|)$ is a seminormed \mathcal{O}_{X_G} -module. For any analytic point x the kernel of $\|\cdot\|_x$ contains $m_{G,x}\mathcal{F}_x$ and hence $\|\cdot\|_x$ is induced from the residue seminorm on the fiber $\mathcal{F}(x) = \mathcal{F}_x/m_{G,x}\mathcal{F}_x$. We call the latter the *fiber seminorm* and denote it $\|\cdot\|_{(x)}$. The completion of $\mathcal{F}(x)$ with respect to $\|\cdot\|_{(x)}$ is a Banach $\mathcal{H}(x)$ -space that will be called the *completed fiber* and denoted $\widehat{\mathcal{F}(x)}$.

3.3.8 Pullbacks

If $f : Y \rightarrow X$ is morphism of Berkovich spaces and \mathcal{F} is a seminormed \mathcal{O}_{X_G} -module, then we define the pullback as $f^*\mathcal{F} = f^{-1}\mathcal{F} \otimes_{f^{-1}\mathcal{O}_{X_G}} \mathcal{O}_{Y_G}$, where both f^{-1} and the tensor product are taken in the sense of seminormed sheaves.

3.3.9 The Case of Invertible Sheaves

For illustration, let us describe G -analytic \mathcal{O}_{X_G} -seminorms on an invertible module \mathcal{F} . Such a seminorm $\|\cdot\|$ is determined by a G -semicontinuous family of seminorms $\|\cdot\|_x$ for $x \in X$. Sending a seminorm to its fiber establishes a bijection between $\mathcal{O}_{X_G,x}$ -seminorms on \mathcal{F}_x and $\kappa_G(x)$ -seminorms on the one-dimensional vector space $\mathcal{F}(x)$. Finally, if s_x is a basis of $\mathcal{F}(x)$, then sending a seminorm to its value on s_x provides a parametrization of $\kappa_G(x)$ -seminorms on $\mathcal{F}(x)$ by numbers $r \in \mathbf{R}_{\geq 0}$.

Lemma 3.3.10. *Assume that X is an analytic space and \mathcal{F} is a free \mathcal{O}_{X_G} -module of rank one with basis s . Then the correspondence $\|\cdot\| \mapsto \|s\|$ establishes a bijection between G -analytic \mathcal{O}_{X_G} -seminorms on \mathcal{F} and upper G -semicontinuous functions $r : X \rightarrow \mathbf{R}_{\geq 0}$. Furthermore, a seminorm is analytic if and only if the function $\|s\|$ is upper semicontinuous.*

Proof. For any seminorm $\|\cdot\|$ the function $\|s\|$ is upper G -semicontinuous and the argument above the lemma shows that it determines a G -analytic seminorm uniquely. Conversely, given a G -semicontinuous function r consider the seminorm $\|\cdot\|_x$ on \mathcal{F}_x such that $\|s\|_x = r(x)$. It is easy to see that the family $\{\|\cdot\|_x\}_{x \in X}$ is G -semicontinuous and hence gives rise to a G -analytic seminorm by Lemma 3.2.8(ii). It remains to show that if $\|s\|$ is upper semicontinuous then the seminorm is analytic. Indeed, a section $t \in \mathcal{F}(U)$ is of the form fs with $f \in \mathcal{O}_{X_G}(U)$, and using that the function $|f|$ is continuous we obtain that $\|t\| = |f| \cdot \|s\|$ is upper semicontinuous. \square

3.3.11 Analytic \mathcal{O}_{X_G} -Seminorms on Coherent Sheaves

For coherent sheaves one can describe analyticity in terms of fiber seminorms of X and X_G .

Lemma 3.3.12. *Assume that X is a good analytic space with a coherent \mathcal{O}_X -module \mathcal{F} and let \mathcal{F}_G denote the associated coherent \mathcal{O}_{X_G} -module on X_G . Let $\|\cdot\|$ be an \mathcal{O}_{X_G} -seminorm on \mathcal{F}_G and for any $x \in X$ endow stalks \mathcal{F}_x and $\mathcal{F}_{G,x}$ and the fibers $\mathcal{F}(x)$ and $\mathcal{F}_G(x)$ with the stalk and the fiber seminorms of $\|\cdot\|$. Then the following conditions are equivalent:*

- (i) $\|\cdot\|$ is analytic.
- (ii) The map $\mathcal{F}_x \rightarrow \mathcal{F}_{G,x}$ is an isometry for any $x \in X$.
- (iii) The map $\mathcal{F}(x) \rightarrow \mathcal{F}_G(x)$ is an isometry for any $x \in X$.
- (iv) The map $\mathcal{F}(x) \otimes_{\kappa(x)} \kappa_G(x) \rightarrow \mathcal{F}_G(x)$ is an isometric isomorphism for any $x \in X$.

Proof. (iii) \iff (iv) Set $K = \kappa(x)$, $K' = \kappa_G(x)$, $V = \mathcal{F}(x)$ and endow $V' = V \otimes_K K'$ with the tensor seminorm. The map $h : V' \rightarrow \mathcal{F}_G(x)$ is an isomorphism since \mathcal{F} is coherent, so (iv) is satisfied if and only if h is an isometry. Note also that the inclusion $K \hookrightarrow K'$ is an isometry, hence $V \hookrightarrow V'$ is an isometry by [32, Lemma 3.1]. In particular, if h is an isometry, then $g : V \rightarrow \mathcal{F}_G(x)$ is an isometry, i.e. (iv) \implies (iii). Conversely, if g is an isometry, then h is an isometry because h is non-expansive and V is dense in V' .

(ii) \iff (iii) This follows from the fact that the seminorm of \mathcal{F}_x is induced from the seminorm of $\mathcal{F}(x)$, and similarly for $\mathcal{F}_{G,x}$ and $\mathcal{F}_G(x)$.

(i) \implies (ii) We should prove that if $x \in X$ and $s \in \mathcal{F}_x$ then the image $s_G \in \mathcal{F}_{G,x}$ of s satisfies $\|s_G\|_x \geq \|s\|$. Choose a neighborhood U of x such that s is defined on U . Since the seminorm is analytic, for any $r > \|s_G\|_x$ there exists a neighborhood $V \subseteq U$ of x such that $\|s\|_y < r$ for any $y \in V$. Since $\|\cdot\|$ is G -analytic, this implies that $\|s\|_V < r$ and hence $\|s\|_x < r$.

(ii) \implies (i) First, we claim that (ii) holds for the restriction of \mathcal{F}_G onto any good domain $U \subseteq X$. Let \mathcal{G} denote the coherent \mathcal{O}_U -module $\mathcal{F}_G|_U$. Choose any $x \in U$ and let $\kappa_U(x)$ be the residue field of x in U , in particular, $\kappa(x) \subseteq \kappa_U(x) \subseteq \kappa_G(x)$. Since (ii) and (iii) are equivalent, we should check that the map $h : \mathcal{G}(x) \rightarrow \mathcal{G}_G(x) = \mathcal{F}_G(x)$ is an isometry. This follows easily from the fact that $\mathcal{G}(x) = \mathcal{F}(x) \otimes_{\kappa(x)} \kappa_U(x)$ as vector spaces, h is non-expansive and $\mathcal{F}(x) \rightarrow \mathcal{F}_G(x)$ is an isometry.

Now, let us prove that $\|\cdot\|$ is analytic. Let $U \subseteq X$ be an affinoid domain and $s \in \mathcal{F}_G(U)$. We claim that the function $\|s\|$ is upper semicontinuous. It suffices to check this for affinoid domains in U hence we can assume that U is affinoid. Fix $x \in U$ and let $r > \|s\|_x$. We showed that $\mathcal{G}_x \rightarrow \mathcal{F}_{G,x}$ is an isometry, where $\mathcal{G} = \mathcal{F}_G|_U$, hence there exists a neighborhood $V \subseteq U$ of x with $\|s\|_V < r$. In particular, $\|s\|_y < r$ for any $y \in V$ and hence $\|s\|$ is upper semicontinuous at x .

It remains to show that $\|\cdot\|$ is G -analytic. If this is not so, then $\|s\|_U > r > \sup_{x \in U} \|s\|_x$ for some choice of U , $s \in \mathcal{A}(U)$ and r . If $U = \cup_i U_i$ is an admissible affinoid covering, then $\|s\|_{U_i} > r$ for some i , and replacing U with U_i we can assume that U is affinoid. By the same argument as above, for any point $x \in U$ there exists an open neighborhood $V_x \subseteq U$ of x with $\|s\|_{V_x} < r$. Since $U = \cup_x V_x$ is an admissible covering, we obtain that $\|s\|_U \leq \sup_x \|s\|_{V_x} < r$. The contradiction concludes the proof. \square

3.3.13 Operations

Finally, let us study when the property of a seminorm to be analytic is preserved by various operations. Certainly, some restrictions should be imposed (see Remark 3.2.9), so we only consider the coherent case, which will be used later.

Lemma 3.3.14. *Analyticity of \mathcal{O}_{X_G} -seminorms on coherent sheaves is preserved under the following operations: quotients, tensor products, symmetric and exterior powers.*

Proof. We consider the case of quotients. The tensor products are dealt with similarly and the other cases follow. So, let $\mathcal{M} \twoheadrightarrow \mathcal{N}$ be a surjection of coherent \mathcal{O}_{X_G} -modules and let $\|\cdot\|_{\mathcal{M}}$ be an analytic \mathcal{O}_{X_G} -seminorm on \mathcal{M} . We should prove that the quotient seminorm $\|\cdot\|_{\mathcal{N}}$ is analytic. By Remark 3.3.4(ii), the question is G -local on X , hence we can assume that X is good.

In the sequel, \mathcal{M} and \mathcal{N} denote the \mathcal{O}_X -sheaves, while \mathcal{M}_G and \mathcal{N}_G denote the associated \mathcal{O}_{X_G} -sheaves. Set $K = \kappa(x)$ and $K' = \kappa_G(x)$, and endow $U = \mathcal{M}(x)$, $U' = \mathcal{M}_G(x)$, $V = \mathcal{N}(x)$ and $V' = \mathcal{N}_G(x)$ with the fiber seminorms of $\|\cdot\|_{\mathcal{M}}$ and $\|\cdot\|_{\mathcal{N}}$. In particular, we have surjections $f : U \twoheadrightarrow V$ and $f' = f \otimes_K K' : U' \twoheadrightarrow V'$ such that the seminorms on the targets are the quotient seminorms.

By the equivalence (i) \iff (iv) in Lemma 3.3.12, the isomorphism $U \otimes_K K' \xrightarrow{\sim} U'$ is an isometry and it suffices to prove that the isomorphism $h : V \otimes_K K' \xrightarrow{\sim} V'$ of seminormed vector spaces is an isometry. By definition, $\|\cdot\|_V$ is the maximal seminorm such that $\|f(u)\|_V \leq \|u\|_U$ for any $u \in U$, hence the seminorm on $V \otimes_K K'$ is the maximal one such that $\|f(u) \otimes a\| \leq |a| \cdot \|u\|_U$ for any $u \in U$ and $a \in K'$. In the same way, the seminorm on V' is the maximal one such that $\|f'(u \otimes a)\| \leq |a| \cdot \|u\|_U$. Since $h(f(u) \otimes a) = f'(u \otimes a)$, the seminorms match and h is an isometry. \square

4 Differentials of Seminormed Rings

4.1 Kähler Seminorms

4.1.1 The Definition

Given a seminormed ring \mathcal{B} and a homomorphism of rings $\phi : A \rightarrow \mathcal{B}$ we equip $\Omega_{\mathcal{B}/A}$ with the *Kähler seminorm* $\|\cdot\|_{\Omega}$ given by the formula

$$\|x\|_{\Omega} = \inf_{x = \sum c_i db_i} \max_i |c_i|_{\mathcal{B}} |b_i|_{\mathcal{B}},$$

where $x \in \Omega_{\mathcal{B}/A}$ and the infimum is over all representations of x as $\sum c_i d_{\mathcal{B}/A}(b_i)$ with $c_i, b_i \in \mathcal{B}$. In fact, we will always work with $\|\cdot\|_{\Omega}$ when ϕ is a non-expansive homomorphism of seminormed rings, but its independence of the seminorm of A will be used.

4.1.2 Universal Properties

Both the Kähler seminorm and the seminormed module $(\Omega_{\mathcal{B}/\mathcal{A}}, \|\cdot\|_\Omega)$ can be characterized by appropriate universal properties.

Lemma 4.1.3. *Assume that $\mathcal{A} \rightarrow \mathcal{B}$ is a non-expansive homomorphism of seminormed rings. Then,*

- (i) $\|\cdot\|_\Omega$ is the maximal \mathcal{B} -seminorm that makes $d_{\mathcal{B}/\mathcal{A}} : \mathcal{B} \rightarrow \Omega_{\mathcal{B}/\mathcal{A}}$ a non-expansive \mathcal{A} -homomorphism.
- (ii) $d_{\mathcal{B}/\mathcal{A}}$ is the universal non-expansive \mathcal{A} -derivation of \mathcal{B} with values in a seminormed \mathcal{B} -module:

$$\mathrm{Hom}_{\mathcal{B}, \mathrm{nonexp}}(\Omega_{\mathcal{B}/\mathcal{A}}, M) \xrightarrow{\sim} \mathrm{Der}_{\mathcal{A}, \mathrm{nonexp}}(\mathcal{B}, M)$$

for any seminormed \mathcal{B} -module M .

Proof. The first claim is obvious. In (ii), we should only prove that any non-expansive \mathcal{A} -derivation $d : \mathcal{B} \rightarrow M$ with values in a seminormed \mathcal{B} -module M factors into a composition of $d_{\mathcal{B}/\mathcal{A}}$ and a non-expansive homomorphism $h : \Omega_{\mathcal{B}/\mathcal{A}} \rightarrow M$. By the usual universal property, we have a unique such factoring with h being a homomorphism of modules, and it remains to show that h is non-expansive. Let $\|\cdot\|'_\Omega$ be defined by $\|x\|'_\Omega = \max(\|x\|_\Omega, \|h(x)\|_M)$. Then it immediately follows that $(\Omega_{\mathcal{B}/\mathcal{A}}, \|\cdot\|'_\Omega)$ is a seminormed \mathcal{B} -module and $d_{\mathcal{B}/\mathcal{A}}$ is non-expansive with respect to $\|\cdot\|'_\Omega$. So, $\|\cdot\|_\Omega = \|\cdot\|'_\Omega$ by (i), and hence h is non-expansive. \square

4.1.4 An Alternative Definition

Similarly to the case of usual rings, one can define the seminormed module $\Omega_{\mathcal{B}/\mathcal{A}}$ in terms of the kernel of $\mathcal{B} \otimes_{\mathcal{A}} \mathcal{B} \rightarrow \mathcal{B}$. The proof reduces to repeating the classical argument (see [27, Sect. 26.C]) and checking that all relevant maps are non-expansive.

Lemma 4.1.5. *Assume that $f : \mathcal{A} \rightarrow \mathcal{B}$ is a non-expansive homomorphism of seminormed rings and I is the kernel of the induced homomorphism of seminormed rings $\mathcal{B} \otimes_{\mathcal{A}} \mathcal{B} \rightarrow \mathcal{B}$. Then the classical isomorphism $\phi : \Omega_{\mathcal{B}/\mathcal{A}} \xrightarrow{\sim} I/I^2$ is an isometry.*

Proof. Recall that ϕ is induced by the derivation $d : \mathcal{B} \rightarrow I/I^2$ given by $db = b \otimes 1 - 1 \otimes b$. By Lemma 4.1.3(ii), it suffices to show that any non-expansive \mathcal{A} -derivation $\partial : \mathcal{B} \rightarrow M$ factors uniquely into a composition of d and a non-expansive homomorphism of \mathcal{B} -modules $h : I/I^2 \rightarrow M$. In fact, a homomorphism h exists and is unique by the classical theory, so we should only check that it is non-expansive.

Let $\mathcal{C} = \mathcal{B} * M$ denote the \mathcal{B} -module $\mathcal{B} \oplus M$ provided with the multiplication $(b, m)(b', m') = (bb', bm' + b'm)$ and the seminorm $|(b, m)|_{\mathcal{C}} = \max(|b|_{\mathcal{B}}, \|m\|_M)$.

The \mathcal{A} -bilinear map $\mathcal{B} \times \mathcal{B} \rightarrow \mathcal{B} * M$ sending (b, b') to $(bb', b\partial(b'))$ is non-expansive, hence it induces a non-expansive homomorphism $\lambda : \mathcal{B} \otimes_{\mathcal{A}} \mathcal{B} \rightarrow \mathcal{B} * M$. By the classical argument, λ vanishes on I^2 and takes I to M , in particular, it induces a non-expansive homomorphism $h : I/I^2 \rightarrow M$. It remains to notice that $h \circ d = \partial$. \square

4.2 Basic Properties of Kähler Seminorms

4.2.1 Fundamental Sequences

First and second fundamental sequences extend to the context of seminormed rings.

Lemma 4.2.2. *Assume that $\mathcal{A} \rightarrow \mathcal{B} \rightarrow \mathcal{C}$ are non-expansive homomorphisms of seminormed rings. Then,*

(i) *The maps of the first fundamental sequence*

$$\Omega_{\mathcal{B}/\mathcal{A}} \otimes_{\mathcal{B}} \mathcal{C} \xrightarrow{g} \Omega_{\mathcal{C}/\mathcal{A}} \xrightarrow{f} \Omega_{\mathcal{C}/\mathcal{B}} \rightarrow 0$$

are non-expansive and f is strictly admissible.

(ii) *If the homomorphism $\phi : \mathcal{B} \rightarrow \mathcal{C}$ is onto and J is its kernel, then the maps of the second fundamental sequence*

$$J/J^2 \rightarrow \Omega_{\mathcal{B}/\mathcal{A}} \otimes_{\mathcal{B}} \mathcal{C} \xrightarrow{g} \Omega_{\mathcal{C}/\mathcal{A}} \rightarrow 0$$

are non-expansive. Furthermore, if $\mathcal{B} \rightarrow \mathcal{C}$ is strictly admissible, then g is strictly admissible.

Proof. This directly follows from the definition of Kähler seminorms. For example, let us check the second assertion in (ii). We should prove that the quotient seminorm does not exceed the Kähler seminorm of $\Omega_{\mathcal{C}/\mathcal{A}}$. The latter is the maximal seminorm satisfying the inequalities $\|dc\| \leq |c|_{\mathcal{C}}$. For any $r > |c|_{\mathcal{C}}$ we can find $b \in \phi^{-1}(c)$ such that $|b|_{\mathcal{B}} < r$. Therefore $\|db\|_{\Omega} < r$ and we obtain that the quotient seminorm of dc does not exceed r . \square

Remark 4.2.3. Even if the map $g : \Omega_{\mathcal{B}/\mathcal{A}} \otimes_{\mathcal{B}} \mathcal{C} \rightarrow \Omega_{\mathcal{C}/\mathcal{A}}$ is an isomorphism, it does not have to be an isometry. Simple examples of this type are obtained when $\mathcal{B} \rightarrow \mathcal{C}$ is an isomorphism but not an isometry.

Corollary 4.2.4. *Let $\mathcal{A} \rightarrow \mathcal{B}$ be a local homomorphism of seminormed local rings and let $k_{\mathcal{A}}$ and $k_{\mathcal{B}}$ be the residue fields provided with the quotient seminorms. If $|m_{\mathcal{B}}|_{\mathcal{B}} = 0$, then the natural map $\Omega_{\mathcal{B}/\mathcal{A}} \rightarrow \Omega_{k_{\mathcal{B}}/k_{\mathcal{A}}}$ is an isometry.*

Proof. By our assumption, the map $\Omega_{\mathcal{B}/\mathcal{A}} \rightarrow \Omega_{\mathcal{B}/\mathcal{A}}/m_{\mathcal{B}}\Omega_{\mathcal{B}/\mathcal{A}}$ is an isometry. Hence the map $\Omega_{\mathcal{B}/\mathcal{A}} \rightarrow \Omega_{k_{\mathcal{B}}/k_{\mathcal{A}}}$ is an isometry by Lemma 4.2.2(ii). Since the map $\mathcal{A} \rightarrow k_{\mathcal{B}}$ factors through $k_{\mathcal{A}}$ there is also an isometry $\Omega_{k_{\mathcal{B}}/k_{\mathcal{A}}} \rightarrow \Omega_{k_{\mathcal{B}}/k_{\mathcal{A}}}$. \square

4.2.5 Base Change

Similarly to the classical case, Kähler differentials of seminormed modules are compatible with base changes.

Lemma 4.2.6. *Assume that $\mathcal{A} \rightarrow \mathcal{B}$ and $\mathcal{A} \rightarrow \mathcal{A}'$ are non-expansive homomorphisms of seminormed rings and set $\mathcal{B}' = \mathcal{B} \otimes_{\mathcal{A}} \mathcal{A}'$. Then $\phi : \Omega_{\mathcal{B}/\mathcal{A}} \otimes_{\mathcal{A}'} \mathcal{B}' \rightarrow \Omega_{\mathcal{B}'/\mathcal{A}'}$ is an isomorphism of seminormed modules.*

Proof. Since ϕ is an isomorphism of modules we should prove that it is an isometry. Let $\| \cdot \|_s$ and $\| \cdot \|_t$ denote the seminorms on the source and on the target, respectively. Recall that $\| \cdot \|_s$ is the maximal seminorm making the bilinear map $\Omega_{\mathcal{B}/\mathcal{A}} \times \mathcal{B}' \rightarrow \Omega_{\mathcal{B}/\mathcal{A}} \otimes_{\mathcal{A}'} \mathcal{B}'$ non-expansive (see Sect. 2.2.9). This fact and Lemma 4.1.3(i) imply that $\| \cdot \|_s$ is the maximal seminorm for which $\|ad(b) \otimes c\|_s \leq |a|_{\mathcal{B}}|b|_{\mathcal{B}}|c|_{\mathcal{B}'}$ for any $a, b \in \mathcal{B}$ and $c \in \mathcal{B}'$. In addition, Lemma 4.1.3(i) implies that $\| \cdot \|_t$ is the maximal seminorm such that $\|xd(y)\|_t \leq |x|_{\mathcal{B}'}|y|_{\mathcal{B}'}$ for any choice of $x, y \in \mathcal{B}'$. Since $ad(b) \otimes c$ goes to $acd(b)$, it follows that any inequality defining $\| \cdot \|_s$ holds also for $\| \cdot \|_t$, and so ϕ is non-expansive.

It remains to check that ϕ^{-1} is non-expansive and for this we will check that any inequality defining $\| \cdot \|_t$ holds for $\| \cdot \|_s$ too. As we noted above, $\| \cdot \|_t$ is defined by the inequalities $\|z\|_t \leq |x|_{\mathcal{B}'}|y|_{\mathcal{B}'}$, where $x, y \in \mathcal{B}'$ and $z = xdy$. Fix $r > |y|_{\mathcal{B}'}$ and find a representation $y = \sum_{i=1}^n b_i \otimes a'_i$ with $b_i \in \mathcal{B}$, $a'_i \in \mathcal{A}'$, and $|b_i|_{\mathcal{B}}|a'_i|_{\mathcal{A}'} \leq r$ for any $1 \leq i \leq n$. Then $\phi^{-1}(z) = \sum_{i=1}^n d(b_i) \otimes a'_i x$ and hence

$$\|\phi^{-1}(z)\|_s \leq \max_i |b_i|_{\mathcal{B}}|a'_i|_{\mathcal{A}'}|x|_{\mathcal{B}'} \leq r|x|_{\mathcal{B}'}$$

Thus $\|\phi^{-1}(xdy)\|_s \leq |y|_{\mathcal{B}'}|x|_{\mathcal{B}'}$ for any $x, y \in \mathcal{B}'$, and we are done. □

4.2.7 Density

Lemma 4.2.8. *If $\phi : \mathcal{B}_0 \rightarrow \mathcal{B}$ is a non-expansive homomorphism of seminormed rings with a dense image, then the Kähler seminorm of $\Omega_{\mathcal{B}/\mathcal{B}_0}$ vanishes.*

Proof. For any $b \in \mathcal{B}$ and $\varepsilon > 0$ there exists $b_0 \in \mathcal{B}_0$ such that $|b - b_0| < \varepsilon$. Hence $\|db\| = \|d(b - b_0)\| \leq \varepsilon$. □

Corollary 4.2.9. *If $\mathcal{A} \rightarrow \mathcal{B}_0 \rightarrow \mathcal{B}$ are homomorphisms of seminormed rings and the image of $\mathcal{B}_0 \rightarrow \mathcal{B}$ is dense, then the homomorphism $\phi : \Omega_{\mathcal{B}_0/\mathcal{A}} \otimes_{\mathcal{B}_0} \mathcal{B} \rightarrow \Omega_{\mathcal{B}/\mathcal{A}}$ has a dense image.*

Proof. By Lemma 4.2.2(ii), $K = \text{Im}(\phi)$ is the kernel of the admissible surjection $\Omega_{\mathcal{B}/\mathcal{A}} \rightarrow \Omega_{\mathcal{B}/\mathcal{B}_0}$. Since the seminorm on $\Omega_{\mathcal{B}/\mathcal{B}_0}$ is trivial by Lemma 4.2.8, it follows that any element of $\Omega_{\mathcal{B}/\mathcal{A}}$ can be approximated by an element of K with any precision, i.e. K is dense. □

4.2.10 Filtered Colimits

We will also need that Kähler seminorms are compatible with filtered colimits.

Lemma 4.2.11. *Assume that \mathcal{A} is a seminormed ring and $\{\mathcal{B}_\lambda\}_{\lambda \in \Lambda}$ is a filtered family of seminormed \mathcal{A} -algebras with non-expansive transition homomorphisms. Let \mathcal{B} be the colimit seminormed algebra, see Sect. 2.2.11. Then $\Omega_{\mathcal{B}/\mathcal{A}} = \text{colim}_\lambda \Omega_{\mathcal{B}_\lambda/\mathcal{A}}$ as seminormed \mathcal{A} -modules.*

Proof. It is a classical result that the modules are isomorphic, so we should compare the seminorms. Let $\|\cdot\|$ be the colimit seminorm on $\Omega_{\mathcal{B}/\mathcal{A}}$. By Lemma 4.1.3(i) and Sect. 2.2.11, $\|\cdot\|$ is the maximal seminorm such that the composed \mathcal{A} -homomorphisms $\mathcal{B}_\lambda \xrightarrow{d_\lambda} \Omega_{\mathcal{B}_\lambda/\mathcal{A}} \rightarrow \Omega_{\mathcal{B}/\mathcal{A}}$ are non-expansive. The latter decompose as $\mathcal{B}_\lambda \xrightarrow{\phi_\lambda} \mathcal{B} \xrightarrow{d} \Omega_{\mathcal{B}/\mathcal{A}}$ and it follows from the definition of $|\cdot|_{\mathcal{B}}$ that $d : \mathcal{B} \rightarrow \Omega_{\mathcal{B}/\mathcal{A}}$ is non-expansive and $\|\cdot\|$ is the maximal seminorm for which this happens. Thus, $\|\cdot\|$ coincides with the Kähler seminorm. \square

4.3 Completed Differentials

4.3.1 The Module $\hat{\Omega}_{\mathcal{B}/\mathcal{A}}$

By $\hat{\Omega}_{\mathcal{B}/\mathcal{A}}$ we denote the completion of $\Omega_{\mathcal{B}/\mathcal{A}}$. It is a Banach \mathcal{B} -module and we call its norm the *Kähler norm*. The corresponding \mathcal{A} -derivation will be denoted $\hat{d}_{\mathcal{B}/\mathcal{A}} : \mathcal{B} \rightarrow \hat{\Omega}_{\mathcal{B}/\mathcal{A}}$.

4.3.2 The Universal Property

Lemma 4.1.3 and the universal property of the completions imply the following result.

Lemma 4.3.3. *If $\phi : \mathcal{A} \rightarrow \mathcal{B}$ is a non-expansive homomorphism of seminormed rings, then $\hat{d}_{\mathcal{B}/\mathcal{A}}$ is the universal non-expansive \mathcal{A} -derivation of \mathcal{B} with values in Banach \mathcal{B} -modules. Namely,*

$$\text{Hom}_{\mathcal{B}, \text{nonexp}}(\hat{\Omega}_{\mathcal{B}/\mathcal{A}}, M) \xrightarrow{\sim} \text{Der}_{\mathcal{A}, \text{nonexp}}(\mathcal{B}, M)$$

for any Banach \mathcal{B} -module M .

Remark 4.3.4. (i) In the case of k -affinoid algebras, Berkovich defines $\hat{\Omega}_{\mathcal{B}/\mathcal{A}}$ in [5, Sect. 3.3] as J/J^2 , where $J = \text{Ker}(\mathcal{B} \rightarrow \mathcal{B} \hat{\otimes}_{\mathcal{A}} \mathcal{B})$. Note that the notation in [5] does not use hat because uncompleted modules of Kähler differentials are never considered there, but we have to distinguish them in our paper. It follows

from [5, Proposition 3.3.1(ii)] and Lemma 4.3.3 that Berkovich’s definition is equivalent to ours. In particular, if $X = \mathcal{M}(\mathcal{B})$ and $S = \mathcal{M}(\mathcal{A})$, then $\Gamma(\Omega_{X/S}) = \widehat{\Omega}_{\mathcal{B}/\mathcal{A}}$.

- (ii) Alternatively, one could deduce the equivalence of the two definitions from Lemma 4.1.5. Even more generally, Lemma 4.1.5 implies that if $\mathcal{A} \rightarrow \mathcal{B}$ is a homomorphism of Banach rings, $J = \text{Ker}(\mathcal{B} \rightarrow \mathcal{B} \widehat{\otimes}_{\mathcal{A}} \mathcal{B})$ and $\overline{J^2}$ is the closure of J^2 in J , then $\widehat{\Omega}_{\mathcal{B}/\mathcal{A}} = \widehat{J/J^2} = J/\overline{J^2}$. In the affinoid case, all ideals are automatically closed, so $\overline{J^2} = J^2$.

5 Kähler Seminorms of Real-Valued Fields

Our next aim is to study Kähler seminorms on the vector spaces $\Omega_{K/A}$, where K is a real-valued field and $\phi : A \rightarrow K$ is a homomorphism of rings. In this case we will use the notation $A^\circ = \phi^{-1}(K^\circ)$. In fact, we are mainly interested in the cases when $A = \mathbf{Z}$ or A is a field, but we will consider an arbitrary A when possible.

5.1 Log Differentials

The aim of this section is to express the Kähler seminorm on $\Omega_{K/A}$ in terms of modules of log differentials, so we start with recalling some basic facts about the latter.

5.1.1 Log Rings

A *log structure* on a ring A is a homomorphism of monoids $\alpha_A : M_A \rightarrow (A, \cdot)$ inducing an isomorphism $M_A^\times \xrightarrow{\sim} A^\times$. The triple (A, M_A, α_A) is called a *log ring*. Usually we will denote it as (A, M_A) and, when this cannot cause to a confusion, denote elements $\alpha_A(m)$ simply by m . Homomorphisms of log rings are defined in the natural way.

5.1.2 Log Differentials

If $(A, M_A) \rightarrow (B, M_B)$ is a homomorphism of log rings, then we denote by $\Omega_{(B, M_B)/(A, M_A)}$ the module of log differentials. Recall that the latter is defined in [23, Sect. 1.7] as the quotient of $\Omega_{B/A} \oplus (B \otimes M_B^{\text{gp}})$ by the relations of the form $(0, 1 \otimes a)$, where $a \in M_A$, and $(d_{B/A}(b), -b \otimes b)$, where $b \in M_B$. The full form of the latter relation is $(d_{B/A}(\alpha_B(b)), -\alpha_B(b) \otimes b)$, and we used our convention to present it in a more compact form. For any $b \in M_B$ we denote by $\delta_{B/A}(b)$ the image

of $1 \otimes b \in B \otimes M_B^{\text{gp}}$ in $\Omega_{(B, M_B)/(A, M_A)}$. It satisfies the equality $b\delta_{B/A}(b) = d_{B/A}(b)$. Intuitively, one may view δb as the log differential $d \log b$ and $\Omega_{(B, M_B)/(A, M_A)}$ is the universal module obtained from $\Omega_{B/A}$ by adjoining the log differentials of the elements of M_B .

5.1.3 The Universal Log Derivative

A log (A, M_A) -derivation of (B, M_B) with values in a B -module N consists of an A -derivation $d : B \rightarrow N$ and a homomorphism $\delta : M_B \rightarrow N$ such that $dm = m\delta m$ for any $m \in M_B$ and $\delta m = 0$ for any m coming from M_A . The B -module of all log derivations is denoted $\text{Der}_{(A, M_A)}((B, M_B), N)$.

Lemma 5.1.4. *If $(A, M_A) \rightarrow (B, M_B)$ is a homomorphism of log rings, then*

$$(d_{B/A} : B \rightarrow \Omega_{(B, M_B)/(A, M_A)}, \delta_{B/A} : M_B \rightarrow \Omega_{(B, M_B)/(A, M_A)})$$

is a universal log (A, M_A) -derivation of (B, M_B) . Namely, for any B -module N the induced homomorphism

$$\text{Hom}_B(\Omega_{(B, M_B)/(A, M_A)}, N) \xrightarrow{\sim} \text{Der}_{(A, M_A)}((B, M_B), N)$$

is an isomorphism.

This fact is proved similarly to its classical analogue for Kähler differentials, so we skip the proof (see also [30, Proposition IV.1.1.6]). As an immediate corollary one obtains the first fundamental sequence (see also [30, Proposition IV.2.3.1]).

Corollary 5.1.5. *Assume that $(A, M_A) \rightarrow (B, M_B) \rightarrow (C, M_C)$ are homomorphisms of log rings. Then the sequence*

$$\Omega_{(B, M_B)/(A, M_A)} \otimes_B C \rightarrow \Omega_{(C, M_C)/(A, M_A)} \rightarrow \Omega_{(C, M_C)/(B, M_B)} \rightarrow 0$$

is exact.

5.1.6 The Integral Log Structure

Given a valued field K and a homomorphism $\phi : B \rightarrow K^\circ$, we will use the notation

$$\Omega_{K^\circ/B}^{\log} = \Omega_{(K^\circ, K^\circ \setminus \{0\})/(B, B \setminus \text{Ker}(\phi))}$$

and $\Omega_{K^\circ}^{\log} = \Omega_{K^\circ/\mathbb{Z}}^{\log}$.

5.1.7 The Kähler Seminorm on $\Omega_{K/A}$

The following result expresses the Kähler seminorm in terms of a module of log differentials.

Theorem 5.1.8. *Assume that K is a real-valued field with a subring A and let $\|\cdot\|_\Omega$ denote the Kähler seminorm of $\Omega_{K/A}$. Then,*

- (i) *Localization induces an isomorphism $\Omega_{K^\circ/A^\circ}^{\log} \otimes_{K^\circ} K = \Omega_{K/A}$. In particular, one can naturally view $(\Omega_{K^\circ/A^\circ}^{\log})_{\text{tf}}$ as a semilattice in $\Omega_{K/A}$.*
- (ii) *$\|\cdot\|_\Omega$ is the maximal K -seminorm whose unit ball contains $(\Omega_{K^\circ/A^\circ}^{\log})_{\text{tf}}$. In particular, $(\Omega_{K^\circ/A^\circ}^{\log})_{\text{tf}}$ is an almost unit ball of $\|\cdot\|_\Omega$.*

Proof. Part (i) follows from the fact that the modules of log differentials are compatible with localizations and the log structure at the generic point of K° is trivial:

$$\Omega_{K^\circ/A^\circ}^{\log} \otimes_{K^\circ} K = \Omega_{(K, K^\times)/(A^\circ, A^\circ \setminus \{0\})}^{\log} = \Omega_{K/A^\circ} = \Omega_{K/A}.$$

Note that $\|\cdot\|_\Omega$ is determined by the inequalities $\|dx\|_\Omega \leq |x|$ with $x \in K$. For any seminorm, $\|dx\| \leq |x|$ if and only $\|d(x^{-1})\| = \|-x^{-2}dx\| \leq |x^{-1}|$, hence already the inequalities $\|\delta x\|_\Omega \leq 1$ with $0 \neq x \in K^\circ$ determine $\|\cdot\|_\Omega$. This proves (ii). \square

In view of Sect. 2.5.3 we obtain the following formula for $\|\cdot\|_\Omega$.

Corollary 5.1.9. *Keep the assumptions of Theorem 5.1.8. Then the restriction of the Kähler seminorm of $\Omega_{K/A}$ to $(\Omega_{K^\circ/A^\circ}^{\log})_{\text{tf}}$ coincides with the adic seminorm of the latter: $\|x\|_\Omega = \|x\|_{\text{adic}}$ for any $x \in (\Omega_{K^\circ/A^\circ}^{\log})_{\text{tf}}$.*

- Remark 5.1.10.*
- (i) We have just seen that the Kähler seminorm on $\Omega_{K/A}$ is tightly related to the torsion free quotient of $\Omega_{K^\circ/A^\circ}^{\log}$. Nevertheless, we will later see (e.g., in Theorem 5.6.4) that the torsion submodules of the modules $\Omega_{K^\circ/A^\circ}^{\log}$ do affect subtle issues related to Kähler seminorms.
 - (ii) It follows from Corollary 5.3.3 below that if A itself is a field and $|A^\times|$ is dense then one can replace $\Omega_{K^\circ/A^\circ}^{\log}$ by Ω_{K°/A° in Theorem 5.1.8 and its corollary. However, in the discretely valued case the discrepancy between the two modules is essential, and one has to stick to logarithmic differentials (e.g, see Sect. 8.3).
 - (iii) Although any serious investigation of ramification theory of valued fields should study the modules Ω_{L°/K° and $\Omega_{L^\circ/K^\circ}^{\log}$ and their torsion, it seems that [18, Chap. 6] is the only source where this topic was systematically explored.

5.2 Main Results on Ω_{L°/K°

In Sects. 5.2–5.5 we will study (logarithmic) differentials of real-valued fields. In fact, all results hold for general valued fields whose height does not exceed one, but we prefer to fix the group of values to keep uniformness of the paper: we still work within the “seminormed framework.”

We start with studying modules Ω_{L°/K° , and we are especially interested in a control on their torsion submodules. To large extent this is based on results of Gabber and Ramero from [18, Chap. 6], where arbitrary valued fields are studied.

5.2.1 Basic Ramification Theory Notation

We say that a finite extension L/K is *unramified* if L°/K° is étale. An algebraic extension L/K is called *unramified* if its finite subextensions are so. As in [18, Sects. 6.2], we denote by K^{sh} the strict henselization of K (in the sense that $(K^{\text{sh}})^\circ$ is the strict henselization of K°); it is the maximal unramified extension of K . By K^t we denote the maximal tame extension of K ; it is the union of all extensions L/K^{sh} whose degree is invertible in \widetilde{K} . For example, if K is trivially valued, then $K^t = K^{\text{sh}}$ is the separable closure of K .

5.2.2 The Cotangent Complex $\mathbf{L}_{L^\circ/K^\circ}$

The module Ω_{L°/K° is the zeroth homology of the cotangent complex $\mathbf{L}_{L^\circ/K^\circ}$. Studying the latter is the central topic of [18, Sects. 6.3, 6.5] and we formulate the main result below. We say that a field extension L/K is *separable* if L is geometrically reduced over K .

Theorem 5.2.3. *Let L/K be an extension of real-valued fields. Then,*

- (i) $H_i(\mathbf{L}_{L^\circ/K^\circ}) = 0$ for $i > 1$.
- (ii) The module $H_1(\mathbf{L}_{L^\circ/K^\circ})$ is torsion free and it vanishes if and only if L/K is separable.
- (iii) The module Ω_{L°/K° is torsion whenever L/K is algebraic and separable.

Proof. (i) and the first assertion of (ii) are proved in [18, Theorem 6.5.12(i)]. Since $H_1(\mathbf{L}_{L^\circ/K^\circ}) \otimes_{L^\circ} L = H_1(\mathbf{L}_{L/K})$ and the latter module vanishes if and only if L/K is separable, we obtain the second assertion of (ii). Finally, $\Omega_{L^\circ/K^\circ} \otimes_{L^\circ} L = \Omega_{L/K}$ and the latter module vanishes whenever L/K is algebraic and separable. \square

Remark 5.2.4. The proving method of [18, Theorem 6.5.12(i)] is as follows: first one explicitly computes $\mathbf{L}_{L^\circ/K^\circ}$ for certain elementary extensions L/K using that in this case L° is a filtered colimit of subrings $K^\circ[x_i]$ (see [18, Propositions 6.3.13 and 6.5.9]), then the general case is deduced via transitivity triangles.

5.2.5 The Six-Term Exact Sequence

In the sequel, we will use the notation $\Upsilon_{L^\circ/K^\circ} = H_1(\mathbf{L}_{L^\circ/K^\circ})$. By [22, Proposition II.2.1.2] any tower of real-valued extensions $F/L/K$ gives rise to a distinguished transitivity triangle of cotangent complexes

$$\mathbf{L}_{L^\circ/K^\circ} \otimes_{L^\circ} F^\circ \rightarrow \mathbf{L}_{F^\circ/K^\circ} \rightarrow \mathbf{L}_{F^\circ/L^\circ} \xrightarrow{+1}$$

and by Theorem 5.2.3 we obtain a six-term exact sequence of homologies

$$0 \rightarrow \Upsilon_{L^\circ/K^\circ} \otimes_{L^\circ} F^\circ \rightarrow \Upsilon_{F^\circ/K^\circ} \rightarrow \Upsilon_{F^\circ/L^\circ} \rightarrow \Omega_{L^\circ/K^\circ} \otimes_{L^\circ} F^\circ \rightarrow \Omega_{F^\circ/K^\circ} \rightarrow \Omega_{F^\circ/L^\circ} \rightarrow 0.$$

5.2.6 Tame Extensions

Cotangent complexes of tame extensions are especially simple.

Lemma 5.2.7. *If L/K is a tame algebraic extension, then the L° -module Ω_{L°/K° is isomorphic to $L^\circ/K^\circ \otimes L^\circ$ and hence $\mathbf{L}_{L^\circ/K^\circ}$ is quasi-isomorphic to the module $L^\circ/K^\circ \otimes L^\circ$ placed in the degree 0.*

Proof. By Theorem 5.2.3 it suffices to prove the first isomorphism. The claim is obvious for unramified extensions since both sides vanish. Using the transitivity triangles we can now replace L and K with L^{sh} and K^{sh} and assume in the sequel that $K = K^{\text{sh}}$. Both sides of the asserted equality commute with filtered unions of extensions, hence we can assume that L/K is finite. Then the lemma reduces to two cases: (a) if the valuation of K is discrete K , then Ω_{L°/K° is cyclic of length $e_{L/K} - 1$, (b) if $|K^\times|$ is dense, then $\Omega_{L^\circ/K^\circ} = 0$. The first claim is classical, so we only check the second one.

Since $K = K^{\text{sh}}$ it follows that L/K breaks into a tower of elementary extensions of the form $F(\pi^{1/l})/F$ with l invertible in \tilde{K} and $|\pi| \notin |K^\times|^l$. Using the transitivity triangles it suffices to consider the case when $L = K(\pi^{1/l})$. Note that L° is the filtered union of the subalgebras $A_i = K^\circ[x_i^{1/l}]$ where $x_i \in \pi K^l$ satisfy $|x_i| < 1$. Since $l \in (L^\circ)^\times$ one easily sees that $\Omega_{A_i/K^\circ} = A_i dx_i / x_i^{l-1} A_i dx_i$, and hence Ω_{L°/K° is the filtered colimit of the modules $M_j = L^\circ dx_i / x_i^{l-1} L^\circ dx_i$ with the following transition maps: if $|x_i| \leq |x_j|$, then $A_i \subseteq A_j$ and the map $M_i \rightarrow M_j$ is given by $dx_i \mapsto \frac{x_j}{x_i} dx_j$. The image of dx_i in M_j vanishes whenever $\frac{|x_i|}{|x_j|} \leq |x_j|^{l-1}$, that is $|x_i| \leq |x_j|^l$. Since $|K^\times|$ is dense, $|x_j|$ can be arbitrarily close to 1 and taking $|x_j| > |x_i|^{1/l}$ we kill the image of dx_i . Thus $\text{colim}_i M_i = 0$ and we are done. \square

5.2.8 Dense Extensions

Using the method outlined in Remark 5.2.4 we will also study $\mathbf{L}_{L^\circ/K^\circ}$ when K is dense in L . It is easy to see that $\widehat{\Omega}_{L^\circ/K^\circ} = 0$ and hence $\Omega_{L^\circ/K^\circ} = 0$ is divisible, but we will need a more precise statement. We say that an L° -module M is a *vector space* if it is an L -module. Equivalently, M is divisible and torsion free.

Lemma 5.2.9. *Assume that L/K is an extension of real-valued fields and K is dense in L . Then both Ω_{L°/K° and $\Upsilon_{L^\circ/K^\circ}$ are L -vector spaces.*

Proof. We start with three classes of elementary extensions, and then the general case will be deduced in three more steps.

Case 1. *Assume that L/K is separable algebraic.* Since K is dense in L , it follows that L lies in the henselization of K , i.e. L°/K° is étale. In this case, the cotangent complex vanishes.

Case 2. *Assume that L/K is purely inseparable of degree p .* In this case, $L = K(x)$ with $a = x^p \in K \setminus K^p$ and there exist $c_i \in K$ such that $\lim_i |x - c_i| = 0$. Clearly, $f(t) = t^p - a$ is the minimal polynomial of x . It is easy to see that L° is the filtered colimit of its subrings $K^\circ[x_i]$, with $x_i = \frac{x - c_i}{\pi_i}$ where $\pi_i \in K$ are such that $\lim_i |\pi_i| = 0$ (same argument as in the proof of [18, 6.3.13(i)]). It follows that $\mathbf{L}_{L^\circ/K^\circ} = \text{colim}_i \mathbf{L}_{K^\circ[x_i]/K^\circ}$. Since $K^\circ[x_i] = K^\circ[t]/I_i$ where $I_i = (f_i)$ and $f_i(t) = t^p - \frac{a - c_i^p}{\pi_i^p}$ is the minimal polynomial of x_i , the homologies of $\mathbf{L}_{K^\circ[x_i]/K^\circ}$ are easily computable (cf. the proof of [18, 6.3.13(iv)]): $\Omega_{K^\circ[x_i]/K^\circ}$ is the invertible module with basis dx_i and $\Upsilon_{K^\circ[x_i]/K^\circ} = I_i/I_i^2$ is the invertible module with basis f_i . Since $dx_i = \frac{dx}{\pi_i}$, we obtain that $\Omega_{L^\circ/K^\circ} = Ldx$.

To describe the map $\Upsilon_{K^\circ[x]/K^\circ} \rightarrow \Upsilon_{K^\circ[x_i]/K^\circ}$ we consider compatible presentations $K^\circ[t] \twoheadrightarrow K^\circ[x]$ and $K^\circ[t_i] \twoheadrightarrow K^\circ[x_i]$, where the connecting maps take t and x to $\pi_i t_i + c_i$ and $\pi_i x_i + c_i$, respectively. The connecting maps induce the map $I/I^2 \rightarrow I_i/I_i^2$ that sends $f = t^p - a$ to $\pi_i^p t_i^p + c_i^p - a = \pi_i^p f_i(t_i)$. This completely determines the filtered family $\Upsilon_{K^\circ[x_i]/K^\circ}$, and since $\lim_i |\pi_i^p| = 0$, its colimit is the one-dimensional L -vector space with basis f .

Case 3. *Assume that $L = K(x)$ is purely transcendental.* In this case we have that $\Upsilon_{L^\circ/K^\circ} \otimes_{L^\circ} L = \Upsilon_{L/K} = 0$, and since $\Upsilon_{L^\circ/K^\circ}$ is torsion free, it actually vanishes. The module Ω_{L°/K° is computed as in Case 2. First, one shows that L° is the filtered colimit of localizations of the K° -algebras $K^\circ[x_i]$ where $x_i = \frac{x - c_i}{\pi_i}$ and $\lim_i |\pi_i| = 0$ (same argument as in [18, 6.3.13(i)] or [18, 6.5.9]). It then follows that Ω_{L°/K° is the colimit of invertible modules $L^\circ dx_i$ and $dx_i = \frac{dx}{\pi_i}$. Thus, $\Omega_{L^\circ/K^\circ} = Ldx$.

Case 4. *Assume that L/K is finite.* We induct on $[L : K]$. If L/K is as in Cases 1 or 2 then we are done. Otherwise, it can be split into a tower $K \subsetneq F \subsetneq L$, and the claim holds true for F/K and L/F by the induction. Therefore, in the six-term exact sequence the terms corresponding to F/K and L/F are vector spaces. It follows easily that the remaining terms are vector spaces too.

Case 5. *Assume that L/K is finitely generated.* This time we induct on $d = \text{tr.deg.}(L/K)$. If $d = 0$, then we are in Case 4. Otherwise choose a

purely transcendental subextension $F = K(x) \subseteq L$ and note that the assertion holds for F/K by Case 3 and for L/F by the induction. The same argument with the six-term sequence completes Case 5.

Case 6. *The general case.* Obviously, L is the filtered colimit of its finitely generated K -subfields L_i and $L^\circ = \text{colim}_i L_i^\circ$. It remains to use that the cotangent complex and the homology are compatible with filtered colimits, and filtered colimits of vector spaces are vector spaces. \square

5.2.10 The Different

We have defined the content of K° -modules in Sect. 2.6. For any separable extension of real-valued fields L/K we define the *different* to be $\delta_{L/K} = \text{cont}((\Omega_{L^\circ/K^\circ})_{\text{tor}})$. In particular, $\delta_{L/K} = 1$ if L is trivially valued. For inseparable extensions we set $\delta_{L/K} = 0$.

Theorem 5.2.11. *Let $F/L/K$ be a tower of algebraic extensions of real-valued fields, then*

- (i) $\delta_{F/K} = \delta_{F/L} \delta_{L/K}$.
- (ii) *If L is not trivially valued and L/K is tame, then $\delta_{L/K} = |K^{\circ\circ}|/|L^{\circ\circ}|$. In particular, if K is not trivially valued, then $\delta_{K^\circ/K} = |K^{\circ\circ}|$.*
- (iii) *If L/K is finite and separable, then $\delta_{L/K} > 0$.*

Proof. If F/K is inseparable, then either F/L or L/K is inseparable and both sides of the equality in (i) vanish. So, assume that F/K is separable. By Theorem 5.2.3 we have a short exact sequence of torsion modules

$$0 \rightarrow \Omega_{L^\circ/K^\circ} \otimes_{L^\circ} F^\circ \rightarrow \Omega_{F^\circ/K^\circ} \rightarrow \Omega_{F^\circ/L^\circ} \rightarrow 0.$$

Hence (i) follows from Theorem 2.6.7 and the fact that for any L° -module M and the F° -module $M' = M \otimes_{L^\circ} F^\circ$ the equality $\text{cont}(M) = \text{cont}(M')$ holds.

If L is not trivially valued, then $\text{cont}(L^{\circ\circ}/K^{\circ\circ}L^\circ) = |K^{\circ\circ}|/|L^{\circ\circ}|$ and hence (ii) follows from Lemma 5.2.7.

Let us prove (iii). Using (i) it suffices to prove that the different of the larger extension K^tL/L does not vanish. Furthermore, $\delta_{K^t/K} > 0$ by (ii), hence it suffices to consider the finite extension K^tL/K^t . Thus, we can assume that $K = K^t$ and then L/K splits to a tower of extensions of degree $p = \text{char}(\widehat{K})$. Using (i) again it suffices to consider the case when $K = K^t$ and $[L : K] = p$. By [18, Proposition 6.3.13] we have that L° is a filtered colimit of subalgebras $K^\circ[a_i x + b_i]$, where $x \in L$ and $a_i, b_i \in K$. We can also assume that $x \in L^\circ$. Note that the numbers $|a_i|$ are bounded by a finite number C because otherwise $x \in \widehat{K}$ and hence \widehat{K} contains the separable extension L/K , which is impossible since K is henselian.

Set $x_i = a_i x + b_i$, then Ω_{L°/K° is generated by the elements dx_i . Let $f(x)$ be the minimal polynomial of x over K . Then dx is annihilated by $f'(x)$ and $f'(x) \neq 0$ since

L/K is separable. It follows that dx_i is annihilated by any π with $|\pi| < |f'(x)|C^{-1}$ and since Ω_{L°/K° is the filtered union of the submodules generated by dx_i , we obtain by Lemma 2.6.6 that $\text{cont}(\Omega_{L^\circ/K^\circ}) \geq |f'(x)|C^{-1} > 0$. \square

Remark 5.2.12. (i) We will not need this, but our definition of the different is compatible with the definition of Gabber and Ramero in the following sense. They introduced in [18, Sect.2.3] a class \mathcal{C} of uniformly almost finitely generated K° -modules M and defined for them the formation of Fitting’s ideals $F_i(M)$. It is not difficult to show that for any uniformly almost finitely generated torsion module M the equality $|F_0(M)| = \text{cont}(M)$ holds. In addition, it is proved in [18, Proposition 6.3.8] that if L/K is a finite separable extension of real-valued fields then Ω_{L°/K° is uniformly finitely generated, and then the different of L/K is defined to be $F_0(\Omega_{L^\circ/K^\circ})$. Thus, in this case $\delta_{L/K}$ equals to the absolute value of the different from [18].

(ii) Classically, one defines the different only for algebraic extensions, but it is a reasonable invariant in the transcendental case too. For example, one can show that if k is not trivially valued and $x \in \mathbf{A}_k^1$ is the maximal point of a disc E then (the non-classical) $\delta_{\mathcal{H}(x)/k}$ is the maximum of $\delta_{\mathcal{H}(z)/k}$ where z runs over rigid points in E . In the algebraic case, the different measures how wild the extension K/k is. This also provides a good intuition for the meaning of the different when K/k is not algebraic.

5.3 Relations Between $\Omega_{L^\circ/K^\circ}^{\log}$ and Ω_{L°/K°

Our next aim is to compare the modules $\Omega_{L^\circ/K^\circ}^{\log}$ and Ω_{L°/K° .

5.3.1 Comparison Homomorphism

Let $\lambda_{K/A}$ denote the natural map $\Omega_{K^\circ/A^\circ} \rightarrow \Omega_{K^\circ/A^\circ}^{\log}$, and let $\lambda_K = \lambda_{K/\mathbf{Z}}$.

Lemma 5.3.2. *For any real-valued field K there is an exact sequence of K° -modules*

$$0 \rightarrow \Omega_{K^\circ} \xrightarrow{\lambda_K} \Omega_{K^\circ}^{\log} \xrightarrow{\rho_K} |K^\times| \otimes \tilde{K} \rightarrow 0,$$

where $\rho_K(x\delta_K y) = |y| \otimes \tilde{x}$ for $x, y \in K^\circ$ and K° acts on $|K^\times| \otimes \tilde{K}$ through \tilde{K} . In particular, if K is not discretely valued (but K may be trivially valued), then λ_K is an almost isomorphism.

Proof. The map λ_L is injective for any valued field L by [18, Corollary 6.4.18(i)]. Furthermore, if L is of a finite height h , then [18, Corollary 6.4.18(ii)] constructs a filtration of $\text{Coker}(\lambda_L)$ of length h with explicitly described quotients. Since

the height of K is one, this filtration degenerates and the cited result yields an isomorphism $\text{Coker}(\lambda_K) \xrightarrow{\sim} |K^\times| \otimes (K^\circ/K^{\circ\circ})$. The explicit description of ρ_K is due to the description of a map ρ in the proof of [18, Proposition 6.4.15] (see the formula for $\tilde{\rho}$ above [18, Claim 6.4.16]). \square

Corollary 5.3.3. *If L/K is an extension of real-valued fields, then $\text{Ker}(\lambda_{L/K})$ is annihilated by any element of $K^{\circ\circ}$ and $\text{Coker}(\lambda_{L/K})$ is annihilated by any element of $L^{\circ\circ}$. In particular, if $|K^\times|$ is dense, then $\lambda_{L/K}$ is an almost isomorphism.*

Proof. Let N be the image of the map $\Omega_{K^\circ}^{\log} \otimes_{K^\circ} L^\circ \rightarrow \Omega_{L^\circ}^{\log}$ and let μ be the composition of $\lambda_K \otimes_{K^\circ} L^\circ$ and the map $\Omega_{K^\circ}^{\log} \otimes_{K^\circ} L^\circ \twoheadrightarrow N$. Applying the snake lemma to the commutative diagram

$$\begin{array}{ccccccc}
 \Omega_{K^\circ} \otimes_{K^\circ} L^\circ & \longrightarrow & \Omega_{L^\circ} & \longrightarrow & \Omega_{L^\circ}/K^\circ & \longrightarrow & 0 \\
 \downarrow \mu & & \downarrow \lambda_L & & \downarrow \lambda_{L/K} & & \\
 0 & \longrightarrow & N & \longrightarrow & \Omega_{L^\circ}^{\log} & \longrightarrow & \Omega_{L^\circ}^{\log}/K^\circ \longrightarrow 0
 \end{array}$$

we obtain an exact sequence

$$0 \rightarrow \text{Ker}(\lambda_{L/K}) \rightarrow \text{Coker}(\mu) \rightarrow \text{Coker}(\lambda_L) \rightarrow \text{Coker}(\lambda_{L/K}) \rightarrow 0. \tag{1}$$

By Lemma 5.3.2, $\text{Coker}(\lambda_L)$ is annihilated by $L^{\circ\circ}$, hence the same is true for $\text{Coker}(\lambda_{L/K})$. Similarly, $\text{Coker}(\lambda_K)$ is annihilated by $K^{\circ\circ}$, hence the same is true for $\text{Coker}(\lambda_K) \otimes_{K^\circ} L^\circ = \text{Coker}(\lambda_K \otimes_{K^\circ} L^\circ)$, for its quotient $\text{Coker}(\mu)$, and for the submodule $\text{Ker}(\lambda_{L/K})$ of $\text{Coker}(\mu)$. \square

Using Lemma 5.3.2 and its corollary one can easily verify the following examples.

- Example 5.3.4.* (i) If K is discretely valued with uniformizer π , then $\text{Coker}(\lambda_K) = \tilde{K}$ with generator $\delta_K(\pi)$. In particular, λ_K is not an almost isomorphism in this case.
- (ii) If $\text{char}(\tilde{K}) = p > 0$ and $|K^\times|$ is p -divisible, then λ_K is an isomorphism.
- (iii) Assume that L/K is a tamely ramified finite extension of discretely valued fields with uniformizers π_K and π_L . Note that $\Omega_{L^\circ/K^\circ}^{\log} = 0$ (in fact, L°/K° is log étale), but $\Omega_{L^\circ/K^\circ} \xrightarrow{\sim} \pi_K L^\circ / \pi_L L^\circ$ is a cyclic module of length $e_{L/K} - 1$. In particular, $\lambda_{L/K}$ is not injective if the tamely ramified extension L/K is not unramified.

5.3.5 The Discrete Valuation Case

If the valuation of K is discrete, then the discrepancy between $\Omega_{L^\circ/K^\circ}^{\log}$ and Ω_{L°/K° can be sensitive. However, the following trick reduces the study of modules $\Omega_{L^\circ/K^\circ}^{\log}$ to the non-discrete case.

Lemma 5.3.6. *Assume that K is a discrete-valued field, $A \rightarrow K$ a homomorphism, and L/K a tamely ramified algebraic extension. Then $\Omega_{K^\circ/A^\circ}^{\log} \otimes_{K^\circ} L^\circ \xrightarrow{\sim} \Omega_{L^\circ/A^\circ}^{\log}$.*

Proof. Since the formation of modules of log differentials is compatible with filtered colimits it suffices to consider the case when L/K is finite. Then the situation is well known: the log structures are fine and the inclusion $K^\circ \hookrightarrow L^\circ$ is log étale. By [31, 1.1(iii)] Olsson’s log cotangent complex $\mathbf{L}_{L^\circ/K^\circ}^{\log}$ is quasi-isomorphic to zero, and the lemma follows by considering the transitivity triangle

$$\mathbf{L}_{K^\circ/A^\circ}^{\log} \otimes_{K^\circ} L^\circ \rightarrow \mathbf{L}_{L^\circ/A^\circ}^{\log} \rightarrow \mathbf{L}_{L^\circ/K^\circ}^{\log} \xrightarrow{+1}$$

of $L^\circ/K^\circ/A^\circ$ (see [31, 1.1(v)]) and the associated sequence of homologies. □

5.4 Main Results on $\Omega_{L^\circ/K^\circ}^{\log}$

One can define the whole log cotangent complex $\mathbf{L}_{L^\circ/K^\circ}^{\log}$ using an approach of Gabber that was elaborated by Olsson in [31, Sect. 8]. The advantage of Gabber’s theory is that it deals with not necessarily fine log rings, as our case is. It seems very probable that analogues of all main results of Sect. 5.2 hold also in the logarithmic setting. We do not explore this here and only study the maps

$$\psi_{L^\circ/K^\circ/A^\circ}^{\log} : \Omega_{K^\circ/A^\circ}^{\log} \otimes_{K^\circ} L^\circ \rightarrow \Omega_{L^\circ/A^\circ}^{\log}$$

by comparing them with the non-logarithmic analogues

$$\psi_{L^\circ/K^\circ/A^\circ} : \Omega_{K^\circ/A^\circ} \otimes_{K^\circ} L^\circ \rightarrow \Omega_{L^\circ/A^\circ}$$

Note that $\psi_{L^\circ/K^\circ/A^\circ}^{\log}$ is obtained from $\psi_{L^\circ/K^\circ/\mathbf{Z}}^{\log}$ by dividing both the source and the target by the image of $\Omega_{A^\circ}^{\log} \otimes_{A^\circ} L^\circ$, in particular, the following result holds.

Lemma 5.4.1. *Keep the above notation then $\text{Coker}(\psi_{L^\circ/K^\circ/\mathbf{Z}}^{\log}) = \text{Coker}(\psi_{L^\circ/K^\circ/A^\circ}^{\log})$ and $\text{Ker}(\psi_{L^\circ/K^\circ/A^\circ}^{\log})$ is a quotient of $\text{Ker}(\psi_{L^\circ/K^\circ/\mathbf{Z}}^{\log})$.*

5.4.2 Tame Extensions

The assertion of Lemma 5.3.6 holds for non-discretely valued fields, but we will only need the following slightly weaker version.

Lemma 5.4.3. *Assume that L/K is a tamely ramified algebraic extension of real-valued fields and $A \rightarrow K$ is a homomorphism. Then the homomorphism $\psi_{L^\circ/K^\circ/A^\circ}^{\log}$ is an almost isomorphism.*

Proof. Lemma 5.3.6 covers the case of a discrete-valued K . The trivially graded case is obvious since the logarithmic structures are trivial. It remains to consider the case when $|K^\times|$ is dense. By Lemma 5.3.2, the maps λ_K and λ_L are almost isomorphisms. Since $\psi_{L^\circ/K^\circ/\mathbb{Z}}$ is an isomorphism by Lemma 5.2.7, we obtain that $\psi_{L^\circ/K^\circ/\mathbb{Z}}^{\log}$ is an almost isomorphism. The assertion for an arbitrary A follows by applying Lemma 5.4.1. \square

5.4.4 Separable Extensions

The logarithmic version of Theorem 5.2.3 would imply that $\psi_{L^\circ/K^\circ/A^\circ}^{\log}$ is injective whenever L/K is separable. Comparison with the non-logarithmic case yields a slightly weaker result:

Lemma 5.4.5. *Assume that L/K is a separable algebraic extension of real-valued fields and $A \rightarrow K$ is a homomorphism. Then the kernel of the map $\psi_{L^\circ/K^\circ/A^\circ}^{\log}$ almost vanishes.*

Proof. Again, Lemma 5.4.1 reduces the general claim to the case of the maps $\psi_{L^\circ/K^\circ/\mathbb{Z}}^{\log}$. Furthermore, using Lemma 5.4.3, it suffices to prove that the kernel of $\psi_{(L')^\circ/(K')^\circ/\mathbb{Z}}^{\log}$ almost vanishes. Both K' and L' are not discrete-valued, hence $\psi_{(L')^\circ/(K')^\circ/\mathbb{Z}}^{\log}$ is almost isomorphic to $\psi_{(L')^\circ/(K')^\circ/\mathbb{Z}}$ by Lemma 5.3.2. The latter map has trivial kernel by Theorem 5.2.3(ii). \square

5.4.6 Dense Extensions

Next we consider the case when K is dense in L .

Lemma 5.4.7. *Assume that L/K is an extension of real-valued fields such that K is dense in L and $A \rightarrow K$ is a homomorphism. Then the L° -modules $\text{Ker}(\psi_{L^\circ/K^\circ/\mathbb{Z}}^{\log})$ and $\text{Coker}(\psi_{L^\circ/K^\circ/A^\circ}^{\log})$ are vector spaces and the L° -module $\text{Ker}(\psi_{L^\circ/K^\circ/A^\circ}^{\log})$ is divisible.*

Proof. Consider the following commutative diagram where $\lambda = \lambda_K \otimes_{K^\circ} L^\circ$ and $\rho = \rho_K \otimes_{K^\circ} L^\circ$:

$$\begin{array}{ccccccc}
 0 & \longrightarrow & \Omega_{K^\circ} \otimes_{K^\circ} L^\circ & \xrightarrow{\lambda} & \Omega_{K^\circ}^{\log} \otimes_{K^\circ} L^\circ & \xrightarrow{\rho} & (|K^\times| \otimes \tilde{K}) \otimes_{K^\circ} L^\circ \longrightarrow 0 \\
 & & \downarrow \psi_{L^\circ/K^\circ/\mathbb{Z}} & & \downarrow \psi_{L^\circ/K^\circ/\mathbb{Z}}^{\log} & & \downarrow \gamma \\
 0 & \longrightarrow & \Omega_{L^\circ} & \xrightarrow{\lambda_L} & \Omega_{L^\circ}^{\log} & \xrightarrow{\rho_L} & |L^\times| \otimes \tilde{L} \longrightarrow 0
 \end{array}$$

The rows are exact by Lemma 5.3.2. Moreover, the formulas for ρ_K and ρ_L provided by Lemma 5.3.2 imply that γ is the natural map. Since $|K^\times| = |L^\times|$ and $\tilde{K} = \tilde{L}$ we obtain that γ is an isomorphism, and hence $\psi_{L^\circ/K^\circ/\mathbb{Z}}$ and $\psi_{L^\circ/K^\circ/\mathbb{Z}}^{\log}$ have the same kernel and cokernel. By Lemma 5.2.9 these L° -modules are vector spaces. The general case follows by applying Lemma 5.4.1. \square

5.4.8 Log Different

Given a separable extension L/K of real-valued fields we define the *log different* of L/K to be $\delta_{L/K}^{\log} = \text{cont}((\Omega_{L^\circ/K^\circ}^{\log})_{\text{tor}})$. For inseparable extensions we set $\delta_{L/K}^{\log} = 0$. The log different is related to the usual different in a very simple way and this allows to easily establish its basic properties:

Theorem 5.4.9. *Let L/K be an algebraic extension of real-valued fields, then*

- (i) *If K is not trivially valued, then $\delta_{L/K}^{\log} = \delta_{L/K} |L^\circ|/|K^\circ|$.*
- (ii) *If F/L is another algebraic extension, then $\delta_{F/K}^{\log} = \delta_{F/L}^{\log} \delta_{L/K}^{\log}$.*
- (iii) *If L/K is tame, then $\delta_{L/K}^{\log} = 1$.*
- (iv) *If L/K is finite and separable, then $\delta_{L/K}^{\log} > 0$.*

Proof. The map $\Omega_{K^\circ}^{\log} \otimes_{K^\circ} (K^t)^\circ \rightarrow \Omega_{(K^t)^\circ}^{\log}$ is an almost isomorphism by Lemma 5.4.3. Using this and the similar fact for L , one obtains an almost isomorphism

$$\Omega_{L^\circ/K^\circ}^{\log} \otimes_{L^\circ} (L^t)^\circ \rightarrow \Omega_{(L^t)^\circ/(K^t)^\circ}^{\log},$$

in particular, $\delta_{L/K}^{\log} = \delta_{L^t/K^t}^{\log}$. By Lemma 5.3.2, $\Omega_{(L^t)^\circ/(K^t)^\circ}^{\log}$ is almost isomorphic to $\Omega_{(L^t)^\circ/(K^t)^\circ}$, hence $\delta_{L^t/K^t}^{\log} = \delta_{L^t/K^t}$. Finally,

$$\delta_{L^t/K^t} = \delta_{L/K} \delta_{L^t/L} / \delta_{K^t/K} = \delta_{L/K} |L^\circ|/|K^\circ|$$

by Theorem 5.2.11, and we obtain (i). Using (i), one immediately deduces the other assertions from their non-logarithmic analogues proved in Theorem 5.2.11. \square

5.5 Almost Tame Extensions and Fields

5.5.1 Almost Unramified Extensions

Let L/K be a separable algebraic extension of real-valued fields. Following Faltings we say that L/K is *almost unramified* if Ω_{L°/K° almost vanishes.

Lemma 5.5.2. *Assume that $F/L/K$ is a tower of separable algebraic extensions of real-valued fields, then*

- (i) L/K is almost unramified if and only if $\delta_{L/K} = 1$.
- (ii) F/K is almost unramified if and only if both F/L and L/K are so.

Proof. Recall that the module Ω_{L°/K° is torsion by Theorem 5.2.3(iii). So, Ω_{L°/K° almost vanishes if and only if its content equals to 1, and we obtain (i). The second claim follows from (i) because the different is multiplicative by Theorem 5.2.11(i). \square

5.5.3 Almost Tame Extensions

In the discrete valuation case, almost unramified is the same as unramified, but in general there even are wildly ramified almost unramified extensions. Thus, being almost unramified is not a good measure of “wildness” of extensions: there are tamely ramified extensions which are not almost unramified, and there are wildly ramified extensions which are almost unramified. Another weird property is that a tame extension L/K is almost unramified if and only if K is not discretely valued.

It is natural to seek for a natural enlargement of the class of almost unramified extensions that includes all tame ones, and this can be achieved very simply: one should pass to logarithmic differentials. So, we say that a separable algebraic extension L/K is *almost tame* if $\Omega_{L^\circ/K^\circ}^{\log}$ almost vanishes. Note that any tame extension is almost tame by Theorem 5.4.9(iii). As earlier, we immediately obtain the following result:

Lemma 5.5.4. *Assume that $F/L/K$ is a tower of separable algebraic extensions of real-valued fields, then*

- (i) L/K is almost tame if and only if $\delta_{L/K}^{\log} = 1$.
- (ii) F/K is almost tame if and only if both F/L and L/K are so.

5.5.5 The Non-discrete Case

In the non-discrete case, there is no difference between almost tame and almost unramified because the kernel and cokernel of $\Omega_{L^\circ/K^\circ} \rightarrow \Omega_{L^\circ/K^\circ}^{\log}$ almost vanish by Corollary 5.3.3.

5.5.6 The Discrete Case

As one should expect, in the discrete-valued case the “almost” version of tameness does not provide anything new.

Lemma 5.5.7. *Assume that L/K is a separable algebraic extension of real-valued fields and the valuation on K is discrete. Then L/K is almost tame if and only if it is tame.*

Proof. Assume first that L/K is finite. Then it is a classical result that $\delta_{L/K} \leq |K^{\circ\circ}|/|L^{\circ\circ}|$ and the equality holds if and only if the extension is tame. Using Theorem 5.4.9(i) we obtain that $\delta_{L/K}^{\log} = 1$ if and only if L/K is tame. The general case follows due to the following two observations: (1) L/K is tame if and only if all its finite subextensions L_i/K are tame, (2) $\delta_{L/K}^{\log}$ is the limit of $\delta_{L_i/K}^{\log}$ because $\Omega_{L^{\circ}/K^{\circ}}^{\log}$ is the filtered colimit of $\Omega_{L_i^{\circ}/K^{\circ}}^{\log}$. \square

5.5.8 The Defectless Case

In fact, one can describe a more general situation where almost tameness reduces to tameness. Recall that a finite extension of real-valued fields L/K is called *defectless* if $e_{L/K}f_{L/K} = [L : K]$, and an algebraic extension is defectless if all its finite subextensions are so.

Lemma 5.5.9. *Assume that L/K is a defectless separable extension of real-valued fields. Then L/K is almost tame if and only if it is tame.*

Proof. The case of a trivially valued K is obvious, and the case of a discretely valued K was established in Lemma 5.5.7, so assume that $|K^{\times}|$ is dense. In this case, L/K is almost tame if and only if it is almost unramified by Sect. 5.5.5, hence the lemma reduces to the following claim: *if $|K^{\times}|$ is dense and L/K is defectless and not tame then L/K is not almost unramified.* We start with establishing special cases of the claim. Lemma 5.5.4(ii) will be used all the time so we will not mention it.

Case 1. L/K is wildly ramified of degree p . If $f_{L/K} = p$ choose $x \in L^{\circ}$ such that $\tilde{x} \notin \tilde{K}$, and if $e_{L/K} = p$ choose $x \in L$ such that $|x| \notin |K|$. In either case, the elements $1, x, \dots, x^{p-1}$ form an orthogonal K -basis of L . We claim that $\delta_{L/K} = r^{1-p}|h'(x)|$, where $h(t)$ is the minimal polynomial of x and $r = |x|$.

If $f_{L/K} = p$, then $L^{\circ} = K^{\circ}[x]$ and hence $\Omega_{L^{\circ}/K^{\circ}} = L^{\circ}dx/L^{\circ}h'(x)dx$. In particular, $\delta_{L/K} = |h'(x)|$, as claimed. If $e_{L/K} = p$, then L° is the filtered union of subrings $A_j = K^{\circ}[x_j]$ with $j \in \mathbf{N}$, where $x_j = \frac{x}{\pi_j}$ and $\pi_j \in K$ are such that $r_j = |\pi_j|$ decrease and tend to r from above. It follows that

$$\Omega_{L^{\circ}/K^{\circ}} = \text{colim}_j \Omega_{A_j/K^{\circ}} = \text{colim}_j L^{\circ}dx_j/L^{\circ}h'_j(x_j)dx_j,$$

where $h_j(t)$ is the minimal polynomial of x_j over K . Note that $h_j(t) = \pi_j^{-p}h(\pi_j t)$. Hence $h'_j(x_j) = \pi_j^{1-p}h'(\pi_j x_j) = \pi_j^{1-p}h'(x)$ and using that $dx_j = \pi_j^{-1}dx$ we obtain

$$\Omega_{L^\circ/K^\circ} = \bigcup_j \pi_j^{-1}K^\circ dx / \bigcup_j \pi_j^{-p}h'(x)K^\circ = K_{r-1}^{\circ\circ} dx / K_{r-p}^{\circ\circ} |h'(x)|,$$

where we set $K_s^{\circ\circ} = \{a \in K \mid |a| > s\}$. Therefore, $\delta_{L/K}$ is as claimed.

It remains to estimate $|h'(x)|$, so let $h(t) = t^p + \sum_{i=0}^{p-1} a_i t^i$. We claim that $|a_i| < r^{p-i}$ for $i > 0$. First, $|a_i| \leq r^{p-i}$ by [8, Proposition 3.2.4/3]. If $e_{L/K} = p$, then $r^{p-i} \notin |K^\times|$ for $0 < i < p$, hence $a_i < r^{p-i}$. If $f_{L/K} = p$ then $\tilde{h}(t)$ is the minimal polynomial of \tilde{x} . Since L/\tilde{K} is inseparable, $\tilde{h}(t)$ is of the form $t^p - c$ and hence $|a_i| < 1 = r$ for $0 < i < p$. Thus,

$$|h'(x)| = \left| px^{p-1} + \sum_{i=1}^{p-1} ia_i x^{i-1} \right| \leq \max \left(|pr^{p-1}|, \max_{0 < i < p} (|a_i r^{i-1}|) \right) < r^{p-1}$$

and we obtain that $\delta_{L/K} < 1$, proving that L/K is not almost tame.

Case 2. L/K is finite, Galois, and totally wildly ramified. In this case, L/K splits into a tower of defectless wildly ramified extensions of degree p . Hence L/K is not almost tame by Case 1.

Case 3. L/K is a composition of a tame extension F/K and a Galois totally wildly ramified extension L/E (i.e. $[L : E] = p^n$). Since L/K is not tame, we have that $n > 0$ and hence L/K is not almost tame by Case 2.

Case 4. L/K is finite. By the standard ramification theory there exists a finite tame extension F/L such that F/K splits into a composition of a tame extension and a Galois totally wildly ramified extension. Since F/L is tame it is defectless and hence F/K is defectless. Thus, F/K is not almost tame by Case 3, and since F/L is almost tame we necessarily have that L/K is not almost tame.

Case 5. *The general case.* By definition, L/K contains a finite non-tame defectless subextension F/K , which is not almost tame by Case 4. Hence L/K is not almost tame. □

5.5.10 Almost Tame Extensions: The Summary

We can summarize a few above results as follows.

Theorem 5.5.11. *Assume that L/K is a separable algebraic extension of real-valued fields, then*

- (i) L/K is almost tame if and only if it is either tame or almost unramified.
- (ii) Assume that L/K is defectless; in particular, this is the case when L/K is tame or $|K^\times|$ is discrete. Then L/K is almost tame if and only if it is tame.
- (iii) If $|K^\times|$ is dense, then L/K is almost tame if and only if L/K is almost unramified.

5.5.12 Deeply Ramified Fields

Assume that the valuation on K is non-trivial. We refer the reader to [18, Sect. 6.6] for the definition and basic properties of deeply ramified fields. Recall that a real-valued field K is called *deeply ramified* if $\Omega_{(K^s)^\circ/K^\circ} = 0$, and by [18, Proposition 6.6.2] this condition can be weakened by requiring that K^s/K is almost unramified (i.e., $\Omega_{(K^s)^\circ/K^\circ}$ almost vanishes). Furthermore, any deeply ramified field is perfect by [18, Proposition 6.6.6(i) \iff (ii)], hence K is deeply ramified if and only if the algebraic closure K^a is almost unramified over K .

5.5.13 Almost Tame Fields

We extend the notion of deeply ramified fields by replacing almost unramified extensions with almost tame ones. Assume that K is a real-valued field, whose valuation can be trivial. We say that K is *almost tame* if the extension K^a/K is almost tame. For example, a trivially valued field is almost tame if and only if it is perfect.

Remark 5.5.14. (i) Giving a name to a special class of valued fields K one often refers either to the property of K over a ground valued field (e.g., over \mathbf{Q}_p) or to the properties all extensions of K satisfy. It is slightly confusing but both approaches are used in the theory of valued fields. On the one hand, deeply ramified extensions are “deeply ramified” over a discretely valued subfield. On the other hand, in the model theory of valued fields, K is called *tame* if the extension K^a/K is tame. Our notion of almost tame fields is an analogue (and generalization) of this classical notion.

(ii) Scholze defines in [33, Sect. 3] perfectoid fields to be complete deeply ramified fields of height one and positive residue characteristic. In a sense, the condition of being perfectoid is a valuation-theoretic version of perfectness. So, it seems natural to extend the class of perfectoid fields by allowing non-complete fields and including perfect trivially valued fields and fields of residue characteristic zero. As we are going to prove, this larger class coincides with the class of almost tame fields. So, the notion “perfectoid” is a reasonable alternative to “almost tame.”

Theorem 5.5.15. *A real-valued field K is almost tame if and only if at least one of the following assertions is true: (i) K is a perfect trivially valued field, (ii) $\text{char}(\widetilde{K}) = 0$, (iii) K is deeply ramified.*

Proof. We can assume that the valuation is non-trivial, as the other case is obvious. If $p = \text{char}(\widetilde{K}) = 0$ then any algebraic extension of K is tamely ramified, so K is almost tame. In the sequel we assume that $p > 0$. If K is discretely valued, then it possesses wildly ramified extensions, so K is not almost tame by Lemma 5.5.7.

If $|K^\times|$ is dense, then there is no difference between almost tame and almost unramified extensions. In particular, in this case K is almost tame if and only if it is deeply ramified. \square

Corollary 5.5.16. *Assume that K is a real-valued field.*

- (i) *If $\text{char}(K) > 0$, then K is almost tame if and only if it is perfect.*
- (ii) *Assume that the valuation is non-trivial. Then K is almost tame if and only if K^s/K is almost tame.*

5.5.17 Separable Extensions of Almost Tame Fields

One can also characterize almost tame fields in terms of separable extensions, at cost of considering transcendental ones.

Theorem 5.5.18. *Let K be a real-valued field. Then K is almost tame if and only if for any separable extension of real-valued fields L/K the module $\Omega_{L^\circ/K^\circ}^{\log}$ is almost torsion free.*

Proof. We start with the direct implication, so assume that K is almost tame. If the valuation of K is trivial, then the module $\Omega_{L^\circ/K^\circ}^{\log} = \text{Coker}(\psi_{L^\circ/K^\circ/\mathbb{Z}}^{\log})$ is torsion free by [18, Corollary 6.5.21], so assume that the valuation is non-trivial. If $p = \text{char}(\tilde{K}) = 0$, then $\Omega_{L^\circ/K^\circ}^{\log}$ is torsion free by [18, Lemma 6.5.16], so assume that $p > 0$. Then $|K^\times|$ is dense and hence $\lambda_{L/K}$ is an almost isomorphism by Corollary 5.3.3. So, it suffices to show that Ω_{L°/K° is torsion free. Recall that K is deeply ramified by Theorem 5.5.15 and hence Ω_{K° is divisible.

Case 1: $\text{char}(K) = p$. Then $\Omega_{L^\circ} = \Omega_{L^\circ/\mathbb{F}_p}$ is torsion free by [18, Claim 6.5.21].

Thus Ω_{L°/K° is the cokernel of the map $\Omega_{K^\circ} \otimes_{K^\circ} L^\circ \rightarrow \Omega_{L^\circ}$ whose source is divisible and target is torsion free. Hence Ω_{L°/K° is torsion free.

Case 2: $\text{char}(K) = 0$. Consider the algebraic closure $F = K^a$ provided with an extension of the valuation. Let E be the composite valued field FL . Then $\Omega_{L^\circ/K^\circ} \otimes_{L^\circ} E^\circ \hookrightarrow \Omega_{E^\circ/K^\circ}$ by [18, Lemma 6.3.32(ii)], hence it suffices to show that Ω_{E°/K° is torsion free. The latter follows from the facts that $\Omega_{F^\circ/K^\circ} = 0$ since K is deeply ramified and Ω_{E°/F° is torsion free by [18, Lemma 6.5.20(i)].

Now, let us prove the inverse implication. The case when the valuation of K is non-trivial follows from Corollary 5.5.16: if the torsion module $\Omega_{L^\circ/K^\circ}^{\log}$ is almost torsion free, then it almost vanishes and hence K^s/K is almost tame. Assume, now, that the valuation is trivial. It suffices to prove that if K is not perfect, say $a \in K \setminus K^p$, then there exists a separable extension L/K such that $\Omega_{L^\circ/K^\circ}^{\log}$ contains an essential torsion element. Let R be the localization of $K[t]$ at the maximal ideal generated by $\pi = t^p - a$; it corresponds to a point $x \in \mathbf{A}_K^1$ with $k(x) = K(a^{1/p})$. Then R is a discrete valuation ring of $K(t)$ with uniformizer π and, since \mathbf{A}_K^1 is a smooth K -curve, $\Omega_{R/K} = Rdt$ is a free module with basis dt . The R -module $\Omega_{R/K}^{\log}$ is

generated over $\Omega_{R/K}$ by $\delta\pi$ subject to the relation $\pi\delta\pi - d\pi = 0$. Since $d\pi = 0$, we obtain that $\Omega_{R/K}^{\log} = R \oplus R/\pi R$ and $\delta\pi$ is a non-trivial torsion element. \square

Remark 5.5.19. It might look surprising that R with the log structure generated by π is not log smooth over K . The reason for this is that the log structure is geometrically “non-reduced” over K because $\pi = (t - a^{1/p})^p$ in $R \otimes_K K(a^{1/p})$.

5.6 Kähler Seminorms and Field Extensions

Assume that L/K is an extension of real-valued fields and $A \rightarrow K$ is a homomorphism of rings. In this section we will apply the theory of real-valued field extensions to compare the Kähler seminorms $\|\cdot\|_{\Omega, K/A}$ and $\|\cdot\|_{\Omega, L/A}$ on $\Omega_{L/A}$ and $\Omega_{K/A}$, respectively.

5.6.1 The Map $\psi_{L/K/A}$

The two seminorms are related by the non-expansive map

$$\psi_{L/K/A} : \Omega_{K/A} \otimes_K L \rightarrow \Omega_{L/A},$$

where the seminorm on the source is the base change of $\|\cdot\|_{\Omega, K/A}$. Naturally, we say that $\|\cdot\|_{\Omega, K/A}$ and $\|\cdot\|_{\Omega, L/A}$ agree if $\psi_{L/K/A}$ is an isometry. For shortness, the seminorms on both sides of $\psi_{L/K/A}$ will be denoted $\|\cdot\|$. This is safe since we consider only one seminorm on each vector space.

To study $\psi_{L/K/A}$ it is useful to consider the following commutative diagram

$$\begin{array}{ccccccc}
 \Omega_{K^\circ/A^\circ}^{\log} \otimes_{K^\circ} L^\circ & \xrightarrow{\psi_{L^\circ/K^\circ/A^\circ}^{\log}} & \Omega_{L^\circ/A^\circ}^{\log} & \longrightarrow & \Omega_{L^\circ/K^\circ}^{\log} & \longrightarrow & 0 \\
 \downarrow \zeta & & \downarrow \varepsilon & & \downarrow \lambda & & \\
 (\Omega_{K^\circ/A^\circ}^{\log} \otimes_{K^\circ} L^\circ)_{\text{tf}} & \xrightarrow{\phi} & (\Omega_{L^\circ/A^\circ}^{\log})_{\text{tf}} & \longrightarrow & \text{Coker}(\phi) & \longrightarrow & 0 \\
 \downarrow \alpha & & \downarrow \beta & & \downarrow \gamma & & \\
 \Omega_{K/A} \otimes_K L & \xrightarrow{\psi_{L/K/A}} & \Omega_{L/A} & \longrightarrow & \Omega_{L/K} & \longrightarrow & 0,
 \end{array} \tag{2}$$

where the top and the bottom rows are the first fundamental sequences. Note that the source and the target of ϕ are the almost unit balls of the two seminorms by Theorem 5.1.8.

Lemma 5.6.2. *Keep the above notation. Then $\psi_{L/K/A}$ is an isometry if and only if $\text{Ker}(\phi)$ is divisible and $\text{Coker}(\phi)$ contains no essential torsion elements.*

Proof. By Corollary 5.1.9, $\psi_{L/K/A}$ is an isometry if and only if ϕ is an isometry with respect to the adic seminorm. It remains to use Lemma 2.3.5(ii). \square

5.6.3 Separable Algebraic Extensions

Note that $\psi_{L/K/A}$ is an isomorphism whenever L/K is separable and algebraic.

Theorem 5.6.4. *Assume that L/K is a separable algebraic extension of real-valued fields and $A \rightarrow K$ is a homomorphism.*

- (i) *If L/K is almost tame, then the isomorphism $\psi_{L/K/A}$ is an isometry.*
- (ii) *Assume that $\Omega_{L^\circ/K^\circ/A^\circ}^{\log}$ is torsion free. Then $\psi_{L/K/A}$ is an isometry if and only if L/K is almost tame. Moreover, the content of the quotient of the unit balls of the two seminorms equals to $\delta_{L/K}^{\log}$.*

Proof. Note that $\Omega_{L^\circ/K^\circ}^{\log}$ is a torsion module because $\Omega_{L^\circ/K^\circ}^{\log} \otimes_{L^\circ} L = \Omega_{L/K} = 0$.

- (i) Being a quotient of $\Omega_{L^\circ/K^\circ}^{\log}$, the module $\text{Coker}(\phi)$ in diagram (2) is almost zero. In addition, $\text{Ker}(\phi) \subseteq \text{Ker}(\psi_{L/K/A}) = 0$. Thus, (i) follows from Lemma 5.6.2.
- (ii) The kernel of $\psi_{L^\circ/K^\circ/A^\circ}^{\log}$ almost vanishes by Lemma 5.4.5, hence its source is almost torsion free. Thus, ε is an isomorphism and ζ is an almost isomorphism, in particular, $\psi_{L^\circ/K^\circ/A^\circ}^{\log}$ is almost isomorphic to ϕ . By Theorem 5.1.8, $\psi_{L^\circ/K^\circ/A^\circ}^{\log}$ is almost isomorphic to the embedding of the unit balls. Therefore, the quotient of the unit balls is almost isomorphic to $\Omega_{L^\circ/K^\circ}^{\log}$ giving rise to equality of the contents. \square

5.6.5 Dense Extensions

Another case that will be very important in the sequel is when K is dense in L , for example, $K = \kappa(x)$ and $L = \mathcal{H}(x)$.

Theorem 5.6.6. *Assume that L/K is an extension of real-valued fields such that K is dense in L , and let $A \rightarrow K$ be a homomorphism of rings. Then the map $\psi_{L/K/A}$ is an isometry with a dense image.*

Proof. Density of the image follows from Corollary 4.2.9. Set $\chi = \psi_{L^\circ/K^\circ/A^\circ}^{\log}$ for shortness. To prove that $\psi_{L/K/A}$ is an isometry we recall that $\text{Coker}(\chi)$ is a vector space and $\text{Ker}(\chi)$ is divisible by Lemma 5.4.7. Let χ_{tor} and χ_{tf} be the maps χ induces between the torsion submodules and the torsion free quotients of its arguments. In particular, χ_{tf} is the map ϕ from diagram (2). The snake lemma yields an exact sequence

$$\text{Ker}(\chi) \xrightarrow{\beta} \text{Ker}(\phi) \rightarrow \text{Coker}(\chi_{\text{tor}}) \xrightarrow{\alpha} \text{Coker}(\chi) \rightarrow \text{Coker}(\phi) \rightarrow 0.$$

Since $\text{Coker}(\chi)$ is torsion free and $\text{Coker}(\chi_{\text{tor}})$ is torsion, $\alpha = 0$ and so $\text{Coker}(\phi) = \text{Coker}(\chi)$ is a vector space. In addition, the torsion free group $\text{Ker}(\phi)$ is an extension of the divisible group $\text{Im}(\beta)$ by the torsion group $\text{Coker}(\chi_{\text{tor}})$. It follows easily that, in fact, $\text{Ker}(\phi) = \text{Im}(\beta)$. Thus, $\text{Ker}(\phi)$ is divisible and hence $\psi_{L/K/A}$ is an isometry by Lemma 5.6.2. □

Corollary 5.6.7. *In the situation of Theorem 5.6.6, the completion of $\hat{\psi}_{L/K/A}$ is an isometric isomorphism. In particular, $\widehat{\Omega}_{K/A} \xrightarrow{\sim} \widehat{\Omega}_{L/A}$.*

5.7 Kähler Seminorms and Monomial Valuations

In this section we study finitely generated extensions L/K such that L°/K° behaves similarly to log smooth extensions, though it does not have to be finitely presented.

5.7.1 Orthonormal Bases of $\Omega_{L/K}$

Note that if t_1, \dots, t_n is a separable transcendence basis of a field extension L/K then dt_1, \dots, dt_n is a basis of $\Omega_{L/K}$. The following result is an immediate corollary of Theorem 5.1.8.

Lemma 5.7.2. *Given an extension of real-valued fields L/K with a separable transcendence basis t_1, \dots, t_n consider the following conditions:*

- (i) $\Omega_{L^\circ/K^\circ}^{\log}$ is a free L° -module with basis $\delta t_1, \dots, \delta t_n$.
- (ii) $\frac{dt_1}{t_1}, \dots, \frac{dt_n}{t_n}$ is an orthonormal basis of $\Omega_{L/K}$.

Then (i) \implies (ii) and the conditions are equivalent whenever $\Omega_{L^\circ/K^\circ}^{\log}$ is torsion free.

5.7.3 Generalized Gauss Valuations

We will use the underline to denote tuples, e.g. $\underline{r} = (r_1, \dots, r_n)$. Assume that K is a real-valued field and $A = K[\underline{t}]$ for $\underline{t} = (t_1, \dots, t_n)$. For any tuple $\underline{r} = (r_1, \dots, r_n) \in \mathbf{R}_{>0}^n$ by $|\cdot|_{\underline{r}}$ we denote the generalized Gauss valuation on A defined by the formula

$$\left| \sum_{i \in \mathbf{N}^n} a_i \underline{t}^i \right|_{\underline{r}} = \max_i (|a_i|_{\underline{r}^i}) = \max_i \left(|a_i| \prod_{j=1}^n r_j^{i_j} \right).$$

The normed ring $(A, |\cdot|_r)$ will be denoted $K[t]_r$. Since $|\cdot|_r$ is multiplicative it extends to a norm on $K(\underline{t})$ that will be denoted by the same letter, and we use the notation $(K(\underline{t}), |\cdot|_r) = K(\underline{t})_r$. The following lemma indicates that $(K(\underline{t})_r)^\circ / K^\circ$ behaves as a log smooth extension.

Lemma 5.7.4. *Assume that K is a real-valued field, $r_1, \dots, r_n > 0$ and $L = K(t_1, \dots, t_n)_r$. Then $\Omega_{L^\circ/K^\circ}^{\log}$ is a free L° -module with basis $\delta t_1, \dots, \delta t_n$.*

Proof. By Lemma 5.1.4 it suffices to show that for any L° -module M and elements m_1, \dots, m_n there exists a unique log K° -derivation $(d, \delta) : (L^\circ, L^\circ \setminus \{0\}) \rightarrow M$ such that $\delta t_i = m_i$. If $u \in K[t]$ satisfies $|u|_r = 1$, then $u \in (L^\circ)^\times$ and $u = \sum_{l \in \mathbb{N}^n} a_l t^l$, where $|a_l t^l| \leq 1$, so we set

$$\delta u = u^{-1} \left(\sum_{l \in \mathbb{N}^n} \sum_{i=1}^n l_i a_l t^l m_i \right).$$

Since $|L^\times| = |K^\times| r_1^{Z_1} \dots r_n^{Z_n}$, an arbitrary element $z \in L^\circ$ is of the form $a t^l u/v$ for $l \in \mathbb{Z}^n$, $a \in K$ and $u, v \in K[t]$ with $|u|_r = |v|_r = 1$, so we set $\delta z = l \delta t + \delta u - \delta v$ and $dz = z \delta z$. It is a direct check that the so-defined (d, δ) is a log K° -derivation. Any other log K° -derivation should satisfy the same formulas for $\delta u, \delta z$ and dz , so uniqueness is clear. □

5.7.5 A Characterization of Gauss Valuations

Under mild technical assumptions, one can also characterize generalized Gauss valuations in terms of $\Omega_{L^\circ/K^\circ}^{\log}$.

Lemma 5.7.6. *Let $L = K(t_1, \dots, t_n)/K$ be a purely transcendental extension of real-valued fields, let $r_i = |t_i|$, let $p = \text{exp.char}(\bar{K})$, and assume that \bar{K} is perfect and $|K^\times|$ is p -divisible. If $\frac{dt_1}{t_1}, \dots, \frac{dt_n}{t_n}$ is an orthonormal basis of $\Omega_{L/K}$, then $L = K(t_1, \dots, t_n)_r$ as a valued field.*

Proof. Assume that, conversely, the valuation $|\cdot|$ on L is not generalized Gauss with respect to \underline{t} . Then the valuation on $K[t]$ is strictly dominated by $|\cdot|_r$, in particular, there exists a polynomial $f(\underline{t}) = \sum_{l \in \mathbb{N}^n} a_l t^l$ in L° such that $|f| < |f|_r = \max_l |a_l| r^l$. Removing all monomials of absolute value strictly smaller than $|f|_r$ we can achieve that the following condition holds: (*) $|f| < |f|_r$ and $|a_l t^l| = |f|_r$ for each $a_l \neq 0$.

Choose f satisfying (*) and of minimal possible degree. We claim that if $p > 1$ then not all monomials are p th powers. Indeed, assume that the claim fails, say $f = \sum_{l \in p\mathbb{N}^n} a_l t^l$. By our assumption on K , for any $a \in K$ there exists $b \in K$ with $|a - b^p| < |a|$. Indeed, since $|K^\times|$ is p -divisible we have that $a = x^p y$ with $|y| = 1$, and using that \bar{K} is perfect we can find $z \in K$ with $\tilde{y} = \bar{z}^p$. Then $b = xz$ is as required. Now, for any a_l fix b_l such that $|b_l^p - a_l| < |a_l|$. Clearly, $\sum_{l \in p\mathbb{N}^n} b_l^p t^l$ satisfies (*) and hence $\sum_{l \in \mathbb{N}^n} b_l t^l$ also satisfies (*), but its degree is smaller than that of f . It follows that there exists $l \in \mathbb{N}^n$ and $1 \leq i \leq n$ such that $a_l \neq 0$ and $|l_i| = 1$.

Since the basis $\frac{dt_1}{t_1}, \dots, \frac{dt_n}{t_n}$ of $\Omega_{L/K}$ is orthonormal, $|f| \geq \|df\|_\Omega$ and

$$df = \sum_{l \in \mathbb{N}^n} \sum_{i=1}^n l_i a_l t^l \frac{dt_i}{t_i},$$

the inequality $|l_i a_l t^l| \leq |f|$ holds for any choice of l and i . However, $|l_i a_l t^l| = |a_l t^l| = |f|_{\underline{t}} > |f|$ for the choice with $a_l \neq 0$ and $|l_i| = 1$, a contradiction. \square

5.7.7 \underline{t} -Monomial Valuations

Let L/K be an extension of real-valued fields and $\underline{t} = (t_1, \dots, t_n)$ a tuple of elements of L . We say the valuation on L is \underline{t} -monomial with respect to K if the induced valuation on $k[\underline{t}]$ is a generalized Gauss valuation. This happens if and only if L contains the valued subfield $K(\underline{t})_{\underline{t}}$ where $r_i = |t_i|$. All results proved in Sect. 5.7 can be summarized as follows.

Theorem 5.7.8. *Assume that L/K is an extension of real-valued fields with a separable transcendence basis t_1, \dots, t_n . Consider the following conditions:*

- (i) *The valuation on L is \underline{t} -monomial and $\Omega_{L^\circ/K(t)^\circ}^{\log} = 0$.*
- (ii) *$\Omega_{L^\circ/K^\circ}^{\log}$ is a free L° -module with basis $\delta t_1, \dots, \delta t_n$.*
- (iii) *$\frac{dt_1}{t_1}, \dots, \frac{dt_n}{t_n}$ is an orthonormal basis of $\Omega_{L/K}$.*
- (iv) *$\|\frac{dt_1}{t_1} \wedge \dots \wedge \frac{dt_n}{t_n}\|_{\Omega^n} = 1$ in $\Omega_{L/K}^n$.*

Then (i) \implies (ii) \implies (iii) \iff (iv). Furthermore, if K is almost tame, then all four conditions are equivalent.

Proof. (i) \implies (ii) Set $F = K(\underline{t})$. By the first fundamental sequence, we have a surjective map $\psi^\circ : \Omega_{F^\circ/K^\circ}^{\log} \otimes_{F^\circ} L^\circ \rightarrow \Omega_{L^\circ/K^\circ}^{\log}$, which becomes the isomorphism

$\psi : \Omega_{F/K} \otimes_F L \xrightarrow{\sim} \Omega_{L/K}$ after tensoring with L . In particular, the kernel of ψ° is torsion. By Lemma 5.7.4, the source of ψ° is a free module with basis $\delta t_1, \dots, \delta t_n$. So, the kernel of ψ° is trivial, and hence ψ° is an isomorphism.

(ii) \implies (iii) This is covered by Lemma 5.7.2.

(iii) \iff (iv) The direct implication is obvious. Conversely, assume that (iii) fails.

Then the lattice M generated by the elements $e_i = \frac{dt_i}{t_i}$ is strictly smaller than $\Omega_{L/K}^\circ$ and hence there exists a lattice M' such that $M \subsetneq M' \subseteq \Omega_{L/K}^\circ$. Choose a basis e'_1, \dots, e'_n of M' then

$$1 \geq \|e'_1 \wedge \dots \wedge e'_n\|_{\Omega^n} = [M' : M] \cdot \|e_1 \wedge \dots \wedge e_n\|_{\Omega^n},$$

and since $[M' : M] > 1$ we obtain that (iv) fails.

Finally, assume that K is almost tame, in particular, $\Omega_{L^\circ/K^\circ}^{\log}$ is torsion free. Then the implication (iii) \implies (ii) follows from Lemma 5.7.2. Assume, now, that (ii)

holds; in particular, ψ° is surjective and hence $\Omega_{L^\circ/F^\circ}^{\log} = 0$. Furthermore, the isomorphism ψ is non-expansive and $\|\frac{dt_i}{t_i}\|_\Omega \leq 1$ in its source. Therefore, ψ is an isometry and $\frac{dt_1}{t_1}, \dots, \frac{dt_n}{t_n}$ is an orthonormal basis of $\Omega_{F/K}$. By Lemma 5.7.6, the valuation on F is generalized Gauss, i.e. L is \underline{t} -monomial. \square

Remark 5.7.9. The assumption that K is almost tame in Theorem 5.7.8 is needed for the implication (iii) \implies (ii) even when $n = 1$ and K is trivially valued. For example, assume that $p = \text{char}(K) > 0$ and $a \in K^\circ$ is such that $\tilde{a} \notin \widetilde{K^p}$ and define the norm on $L = K(t)$ so that $|t^p - a| = r$ for some $r \in (0, 1)$ and the norm is $(t^p - a)$ -monomial. Then one can check that $\Omega_{L^\circ/K^\circ}^{\log}$ is a direct sum of $L^\circ \delta t$ and an essential torsion submodule generated by $\delta(t^p - a)$, and hence $\frac{dt}{t}$ is of norm one. Note that $\Omega_{L^\circ/K^\circ}^{\log}$ is not free in this case, although its torsion free quotient is free with basis δt . In the particular case of the trivial valuation, the same example in a non-complete setting was already used in the end of proof of Theorem 5.5.18.

6 Metrization of $\Omega_{X/S}$

Throughout Sect. 6, $f : X \rightarrow S$ denotes a morphism of k -analytic spaces. In the case of sheaves, $|\cdot|$ will always denote spectral seminorms on the structure sheaves and their stalks, and $\|\cdot\|$ will always denote Kähler seminorms (defined below) on sheaves of pluriforms and their stalks.

6.1 Kähler Seminorm on $\Omega_{X/S}$

In this section we introduce a seminorm $\|\cdot\| = \|\cdot\|_{\Omega_{X/S}}$ on $\Omega_{X/S}$. We will mention Ω and X/S in the notation only when a confusion is possible.

6.1.1 The Definition

The construction is straightforward: we simply sheafify the presheaf of Kähler seminorms on affinoid algebras. Let \mathcal{C} be the full subcategory of X_G whose objects are affinoid domains $V = \mathcal{M}(\mathcal{B})$ in X such that $f(V)$ is contained in an affinoid domain $U = \mathcal{M}(\mathcal{A}) \subseteq S$. Note that $\Omega_{X/S}(V) = \widehat{\Omega}_{\mathcal{B}/\mathcal{A}}$, in particular, $\widehat{\Omega}_{\mathcal{B}/\mathcal{A}}$ is independent of the choice of U . Furthermore, the Kähler seminorm on $\widehat{\Omega}_{\mathcal{B}/\mathcal{A}}$ is the quotient of the Kähler seminorm on $\widehat{\Omega}_{\mathcal{A}/k}$ by Lemma 4.2.2(i), hence it is independent of U too and we can denote it $\|\cdot\|'_{\mathcal{B}/S}$. This construction produces a locally bounded pre-quasi-norm $\|\cdot\|'$ on the restriction of $\Omega_{X/S}$ to \mathcal{C} . Its sheafification is a seminorm on $\Omega_{X/S}|_{\mathcal{C}}$, and we can extend this seminorm to a seminorm $\|\cdot\| = \|\cdot\|_{\Omega_{X/S}}$ on the whole $\Omega_{X/S}$ since \mathcal{C} is cofinal in X_G . We call $\|\cdot\|_{\Omega_{X/S}}$ the *Kähler seminorm* of X/S .

Remark 6.1.2. (i) By definition, if V is a domain in X and $\phi \in \Omega_{X/S}(V)$, then $\|\phi\|_V = \inf \max_i \|\phi\|'_{\mathcal{A}_i/S}$, where the infimum is over admissible coverings $V = \cup_i V_i$ with $V_i = \mathcal{M}(\mathcal{A}_i)$ in \mathcal{C} .

(ii) We do not study the question whether $\|\cdot\|_V = \|\cdot\|'_V$ for any V in \mathcal{C} (i.e., whether $\|\cdot\|'$ is already a seminorm on $\Omega_{X/S|C}$). This is not essential for our needs because all results about $\|\cdot\|$ will be proved using stalks.

6.1.3 The Universal Property

The universal property satisfied by Kähler seminorms of seminormed rings, see Lemma 4.1.3(i), and the universal property of sheafification, see Lemma 3.1.5, imply the following characterization of $\|\cdot\|$.

Lemma 6.1.4. *The Kähler seminorm $\|\cdot\|_{\Omega_{X/S}}$ is the maximal seminorm on $\Omega_{X/S}$ that makes the differential $d : \mathcal{O}_X \rightarrow \Omega_{X/S}$ a non-expansive map.*

6.1.5 Unit Balls

There is an alternative approach to Kähler seminorms via unit balls. It is more elementary but less robust and we will not use it. For the sake of completeness we outline it in the following remark leaving some simple verifications to the interested reader.

Remark 6.1.6. (i) If $|k^\times|$ is dense, then Lemma 6.1.4 provides a simple way to describe the unit ball $\Omega_{X/S}^\diamond$ of $\|\cdot\|$, and we obtain another way to define $\|\cdot\|$. This involves the sheaf-theoretic extension of the terminology of Sect. 2.4. Let $\mathcal{B}_{X/S} = \mathcal{O}_{X_G}^\circ d\mathcal{O}_{X_G}^\circ$ denote the subsheaf of $\mathcal{O}_{X_G}^\circ$ -modules of $\Omega_{X/S}$ generated by $d\mathcal{O}_{X_G}^\circ$. Lemma 6.1.4 implies that $\|\cdot\|$ is the maximal seminorm whose unit ball contains $d\mathcal{O}_{X_G}^\circ$, hence $\mathcal{B}_{X/S}$ is an almost unit ball of $\|\cdot\|$ and the unit ball $\Omega_{X/S}^\diamond$ is the almost isomorphic envelope of $\mathcal{B}_{X/S}$.

(ii) Let k be arbitrary. Already for the unit disc $E = \mathcal{M}(k\{t\})$, the inclusion $i : \mathcal{B}_{E/k} \hookrightarrow \Omega_{E/k}^\diamond$ is usually not an isomorphism. For example, if x is the maximal point of a disc around zero of radius $r \notin |k^\times|^\mathbb{Q}$ then $\frac{dt}{t}$ is contained in $\Omega_{E/k,x}^\diamond$ but not in $\mathcal{B}_{E/k,x}$. Moreover, if $|k^\times|$ is discrete, then i is not even an almost isomorphism. This example suggests that one can improve the situation by adding logarithmic differentials of units, so we set

$$\mathcal{B}_{X/S}^{\text{log}} = \mathcal{B}_{X/S} + \mathcal{O}_{X_G}^\circ \delta \mathcal{O}_{X/S}^\times,$$

where $\delta : \mathcal{O}_{X/S}^\times \rightarrow \Omega_{X/S}$ is the logarithmic differential. It is easy to see that $\mathcal{B}_{X/S}^{\text{log}} \subseteq \Omega_{X/S}^\diamond$ and the inclusion is an equality for E/k . It is an interesting question whether $\mathcal{B}_{X/S}^{\text{log}} = \Omega_{X/S}^\diamond$ in general.

6.1.7 The Stalks and the Fibers

Our next aim is to study local behavior of the Kähler seminorm $\| \cdot \|$ at points of X . Given $x \in X$ with $s = f(x)$, fix an affinoid domain $U = \mathcal{M}(\mathcal{A})$ containing s and let $\{V_\lambda = \mathcal{M}(\mathcal{B}_\lambda)\}_\lambda$ be the family of affinoid domains in X such that $x \in V_\lambda$ and $f(V_\lambda) \subseteq U$. Provide $\Omega_{X/S,x}$ with the stalk seminorm $\| \cdot \|_x$, then we saw in Sect. 6.1.1 that $\Omega_{X/S,x}$ is the filtered colimit of the seminormed \mathcal{A} -modules $\widehat{\Omega}_{\mathcal{B}_\lambda/\mathcal{A}}$ (see Sect. 2.2.11). In fact, $\Omega_{X/S,x}$ is an (uncompleted) filtered colimit of completed modules of differentials, so it can be informally thought of as a partial completion of $\Omega_{\mathcal{O}_x/\mathcal{O}_s}$, where we set $\mathcal{O}_x = \mathcal{O}_{X_G,x}$ and $\mathcal{O}_s = \mathcal{O}_{S_G,s}$ for shortness. Consider the following commutative diagram of seminormed modules

$$\begin{array}{ccccc}
 \Omega_{\mathcal{B}_\lambda/\mathcal{A}} & \xrightarrow{\alpha_\lambda} & \widehat{\Omega}_{\mathcal{B}_\lambda/\mathcal{A}} & & \\
 \downarrow \beta_\lambda & & \downarrow \gamma_\lambda & \searrow \psi_\lambda & \\
 \Omega_{\mathcal{O}_x/\mathcal{O}_s} & \xrightarrow{\alpha_x} & \Omega_{X/S,x} & \xrightarrow{\psi_x} & \widehat{\Omega}_{\mathcal{H}(x)/\mathcal{H}(s)},
 \end{array}$$

where ψ_x is the non-expansive \mathcal{A} -homomorphism induced by the \mathcal{A} -homomorphisms ψ_λ via the universal property of colimits.

Theorem 6.1.8. *Keep the above notation. Then α_x and ψ_x are isometries with dense images. In particular, $\widehat{\Omega}_{\mathcal{H}(x)/\mathcal{H}(s)}$ is the completion of both $\Omega_{X/S,x}$ and $\Omega_{\mathcal{O}_x/\mathcal{O}_s}$.*

Proof. By Lemma 4.2.11, $\Omega_{\mathcal{O}_x/\mathcal{O}_s} = \text{colim}_\lambda \Omega_{\mathcal{B}_\lambda/\mathcal{A}}$ as seminormed modules. In particular, α_x is the colimit of isometries with dense images α_λ , and hence α_x itself is an isometry with a dense image. Therefore, it suffices to show that the map $\Omega_{\mathcal{O}_x/\mathcal{O}_s} \rightarrow \widehat{\Omega}_{\mathcal{H}(x)/\mathcal{H}(s)}$ is the completion homomorphism. Indeed, the surjection $\Omega_{\mathcal{O}_x/\mathcal{O}_s} \rightarrow \Omega_{\kappa_G(x)/\kappa_G(s)}$ is an isometry by Corollary 4.2.4 and $\widehat{\Omega}_{\kappa_G(x)/\kappa_G(s)} = \widehat{\Omega}_{\mathcal{H}(x)/\mathcal{H}(s)}$ by Corollary 5.6.7. \square

Corollary 6.1.9. *In the situation of Theorem 6.1.8, ψ_x identifies the completed fiber $\widehat{\Omega_{X/S}(x)}$ with $\widehat{\Omega}_{\mathcal{H}(x)/\mathcal{H}(s)}$.*

Remark 6.1.10. (i) Corollary 6.1.9 provides an alternative way to define $\| \cdot \|_x$.

Instead of the colimit definition in Sect. 6.1.7, one can simply induce $\| \cdot \|_x$ from $\widehat{\Omega}_{\mathcal{H}(x)/\mathcal{H}(s)}$ by the rule $\| \phi \|_x = \| \psi_x(\phi) \|_{\widehat{\Omega}_{\mathcal{H}(x)/\mathcal{H}(s)}}$ for $\phi \in \Omega_{X/S,x}$.

(ii) Even more importantly, the corollary provides a convenient way to compute the values of $\| \cdot \|_x$ since the finite-dimensional normed vector spaces $\widehat{\Omega}_{\mathcal{H}(x)/\mathcal{H}(s)}$ are often pretty explicit. For comparison, we note that it is not clear how to describe the fiber $\Omega_{X/S}(x)$ in terms of $\kappa_G(x)/\kappa_G(s)$. One can only say that $\Omega_{X/S}(x)$ is a finite-dimensional quotient, in fact, a partial completion of the huge vector space $\Omega_{\kappa_G(x)/\kappa_G(s)}$. However, the seminorm of $\Omega_{X/S}(x)$ can still have a non-trivial kernel.

6.1.11 Kähler Seminorms on Pluriforms

By the constructions of Sect. 3.1.7, $\|\cdot\|_\Omega$ induces seminorms on the sheaves obtained from $\Omega_{X/S}$ by tensor products, symmetric powers and exterior powers. In particular, it induces a canonical seminorm $\|\cdot\|_{(\Omega_X^1)^{\otimes m}}$ on the sheaf of pluriforms $(\Omega_X^1)^{\otimes m}$, that will be called Kähler seminorm too.

Of particular interest will be the situation when $X \rightarrow S$ is quasi-smooth of relative dimension n . Then the relative canonical sheaf $\omega_{X/S} = \bigwedge^n \Omega_{X/S}$ is invertible, as well as the relative pluricanonical sheaves $\omega_{X/S}^{\otimes m}$. The corresponding seminorms will be denoted $\|\cdot\|_\omega$ and $\|\cdot\|_{\omega^{\otimes m}}$.

6.1.12 Analyticity of the Seminorms

We have already used Corollary 5.6.7 when studying the stalk seminorms of $\|\cdot\|$. As another application, let us show that all seminorms we have constructed are analytic.

Theorem 6.1.13. *The seminorm $\|\cdot\|_\Omega$ and the induced seminorms on the sheaves obtained from $\Omega_{X/S}$ by tensor products, symmetric powers and exterior powers are analytic.*

Proof. By Remark 3.3.4(ii), analyticity is G -local, hence it suffices to consider the case when X is affinoid. In the sequel $\Omega_{X/S}$ denotes the \mathcal{O}_X -module and the \mathcal{O}_{X_G} -module will be denoted Ω_{X_G/S_G} . By Lemma 3.3.12(i) \iff (iii) we should prove that for each $x \in X$ the map $h : \Omega_{X/S,x} \rightarrow \Omega_{X_G/S_G,x}$ is an isometry with respect to the stalks of $\|\cdot\|_\Omega$. It suffices to check that the completion of h is an isometry and we know by Theorem 6.1.8 that the completion of $\Omega_{X_G/S_G,z}$ is $\widehat{\Omega}_{\mathcal{H}(x)/\mathcal{H}(s)}$. So, it suffices to check that $\widehat{\Omega}_{\mathcal{H}(x)/\mathcal{H}(s)}$ is also the completion of $\Omega_{X/S,x}$, and for this we will copy the argument from the proof of Theorem 6.1.8 but with sheaves in the usual topology.

Set $s = f(x)$, $\mathcal{O}_x = \mathcal{O}_{X,x}$ and $\mathcal{O}_s = \mathcal{O}_{S,s}$, and provide $\Omega_{\kappa(x)/\kappa(s)}$ with the Kähler seminorm. We claim that $\widehat{\Omega}_{\mathcal{H}(x)/\mathcal{H}(s)}$ is the completion of $\Omega_{\mathcal{O}_x/\mathcal{O}_s}$. Indeed, the surjection $\Omega_{\mathcal{O}_x/\mathcal{O}_s} \rightarrow \Omega_{\kappa(x)/\kappa(s)}$ is an isometry by Corollary 4.2.4 and $\widehat{\Omega}_{\kappa(x)/\kappa(s)} = \widehat{\Omega}_{\mathcal{H}(x)/\mathcal{H}(s)}$ by Corollary 5.6.7. □

6.2 Examples

In this section, we compute $\|\cdot\|_\Omega$ and its completed fibers in a few basic cases. We try to choose simple examples that illustrate the general situation. In particular, we will see that $\|\cdot\|_\Omega$ discovers a rather subtle behavior even in the one-dimensional case.

6.2.1 The Case of a Disc

Assume that $k = k^a$ and $X = \mathcal{M}(k\{T\})$ is the unit disc. Then Ω_X is a free sheaf with basis dT , so we can identify it with \mathcal{O}_X by sending dT to 1. Let $r(x)$ be the radius function, i.e. $r(x)$ is the infimum of radii of subdiscs of X containing x . We claim that $\|dT\|_x = r(x)$ for any $x \in X$. If x is contained in a disc of radius s with center at a , then $|T - a|_x \leq s$ and hence $\|dT\|_x = \|d(T - a)\|_x \leq s$. The function $\|dT\| : X \rightarrow \mathbf{R}_{\geq 0}$ is semicontinuous by Theorem 6.1.13, therefore it suffices to check that $\|dT\|_x \geq r(x)$ at a type 2 or 3 point x . In such case, replacing T by a suitable $T - a$ with $a \in k$ we can achieve that $|T|_x = r(x)$ and the valuation on $\mathcal{H}(x)$ is T -monomial. But then $\|\frac{dT}{T}\|_{\Omega, \mathcal{H}(x)/k} = 1$ by Theorem 5.7.8(i) \implies (iii), and so $\|dT\|_x = |T|_x = r(x)$.

The formula for $\|dT\|$ implies that its maximality locus consists of a single point, the maximal point of X . Another consequence is that $\|\cdot\|$ is the seminorm corresponding to $r(x)$ in the sense of Lemma 3.3.10.

Remark 6.2.2. The radius function $r(x)$ is upper semicontinuous but not continuous, and this is a typical behavior for functions of the form $\|\phi\|$. This indicates that the Kähler seminorm on $\Omega_{X/S}$ is very different from the spectral seminorm on \mathcal{O}_{X_G} even when $\Omega_{X/S}$ is invertible. For example, if X is a curve and $f \in \Gamma(\mathcal{O}_{X_G})$ is a global function, then $|f|$ is locally constant outside of a finite graph. On the other hand, if $\phi \in \Gamma(\Omega_X)$ then for any type 2 point x the value of $\|\phi\|$ decreases in almost all directions leading from x . This property is tightly related to the fact that the maximality locus of such ϕ is a finite graph. Also, this indicates that the unit balls $\Omega_{X/S}^\circ$ are usually huge $\mathcal{O}_{X_G}^\circ$ -modules.

6.2.3 Rigid Points: Perfect Ground Field

Assume that k is perfect. Let $x \in X$ be any point with $\mathcal{H}(x) \subseteq \widehat{k^a}$, for example, a rigid point (i.e., a point with $[\mathcal{H}(x) : k] < \infty$), or a type 1 point on a curve. Note that $\widehat{\Omega}_{\mathcal{H}(x)/k} = 0$, since $k^a \cap \mathcal{H}(x)$ is dense in $\mathcal{H}(x)$ by Ax-Sen–Tate theorem, see [2]. Therefore, any differential form ϕ satisfies $\|\phi\|_{\Omega, x} = 0$ by Corollary 6.1.9.

6.2.4 Rigid Points: Non-perfect Ground Field

If k is not perfect, then the situation is different. For example, assume that x is a rigid point with $l = \mathcal{H}(x)$ inseparable over k ; a disc contains plenty of such points, e.g. the points given by $T^p - a = 0$ for $a \in k \setminus k^p$. One can easily give examples when Kähler seminorm on the finite dimensional vector space $\Omega_{l/k}$ is a norm and hence $\widehat{\Omega}_{l/k} = \Omega_{l/k} \neq 0$. (Probably, this is always the case since k is complete.) On the other hand, $\widehat{\Omega}_{l/k}$ is the completion of $\Omega_{X/k, x}$, so the seminorm on the latter does not vanish.

Let us outline a concrete particular case. Assume that $l = k(\alpha)$ is a purely inseparable extension of k of degree p ; in particular, $\Omega_{l/k} = l d\alpha$. Then $r = \inf_{c \in k} |\alpha - c|$ is positive since k is complete. Since $d\alpha = d(\alpha - c)$, we obviously have that $\|d\alpha\|_{\Omega, l/k} \leq r$. One can check straightforwardly that the choice of $\|d\alpha\| = r$ makes the map $d_{l/k}$ non-expansive, so, in fact, $\|d\alpha\|_{\Omega, l/k} = r$. In particular, if T is the coordinate on $X = \mathbf{A}_k^1$ and $x \in X$ is the type 1 point given by $T^p = a$, then $\|dT\|_x = s^{1/p}$, where $s = \inf_{c \in k} |a - c^p|$.

6.2.5 Disc Over Non-perfect Field

Assume that $\text{char}(k) = p > 0$ and \tilde{k} is not perfect and let $X = \mathcal{M}(k\{T\})$. Choose any $a \in k^\circ$ such that $\tilde{a} \notin \tilde{k}^p$ and let x be the rigid point given by $T^p = a$. Then $\|dT\|_x = 1$ by Sect. 6.2.4, and we claim that, more generally, $\|dT\| = 1$ on the whole line connecting x with the maximal point of X . In particular, the maximum locus of $\|dT\|$ contains a huge subgraph, whose combinatorial cardinality (i.e., the cardinality of the set $V \cup E$ of vertices and edges) is easily seen to be equal to the cardinality of k (it is infinite since $\tilde{k} \neq \tilde{k}^p$).

To verify our claim, let q be the maximal point of the disc around x given by $|T^p - a| \leq r$. Set $S = T^p - a$ and $L = \widehat{k(S)} \subseteq \mathcal{H}(q)$, then the valuation on L is S -monomial and hence $\inf_{c \in L} |S + a - c^p| = 1$. Since $T = (a + S)^{1/p}$ we obtain by Sect. 6.2.4 that $\|dT\|_{\Omega, \mathcal{H}(q)/L} = 1$, and hence $\|dT\|_{\Omega, \mathcal{H}(q)/k} \geq 1$. The opposite inequality is obvious, so $\|dT\|_q = \|dT\|_{\Omega, \mathcal{H}(q)/k} = 1$.

Remark 6.2.6. In the mixed characteristic case the situation is even weirder. On the one hand, the Kähler seminorm vanishes at rigid points, but on the other hand, if \tilde{k} is not perfect, say $\tilde{a} \notin \tilde{k}^p$, and I is the interval connecting $x = (T^p - a)$ with the maximal point q then a similar argument shows that $\|dT\| = 1$ on some neighborhood of q in I . In fact, the maximality locus of $\|dT\|$ is a huge tree with root q but its leaves are points of type 2.

6.3 Kähler Seminorms and Base Changes

6.3.1 Domination

For general base changes, Kähler seminorms are related as follows.

Lemma 6.3.2. *Assume that $X \rightarrow S$ and $S' \rightarrow S$ are morphisms of k -analytic spaces and $X' = X \times_S S'$. Let $\phi \in \Gamma(\Omega_{X/S})$ and let $\phi' \in \Gamma(\Omega_{X'/S'})$ be the pullback of ϕ . Then for any point $x' \in X'$ with image $x \in X$ one has that $\|\phi'(x')\| \leq \|\phi(x)\|$. In other words, the pullback of $\|\cdot\|_{\Omega, X/S}$ (see Sect. 3.3.8) dominates $\|\cdot\|_{\Omega, X'/S'}$.*

Proof. Let $s \in S$ and $s' \in S'$ be the images of x' . Unrolling the definitions of the Kähler seminorms and the pullback operation we see that the assertion of the lemma

reduces to the claim that the map $\widehat{\Omega}_{\mathcal{H}(x)/\mathcal{H}(s)} \otimes_{\mathcal{H}(x)} \mathcal{H}(x') \rightarrow \widehat{\Omega}_{\mathcal{H}(x')/\mathcal{H}(x)}$ is non-expansive. The latter follows straightforwardly from the definition of the seminorm on the source: similarly to the argument in Lemma 4.2.6, all inequalities defining the seminorm of the source hold in the target. \square

Remark 6.3.3. (i) The domination of seminorms from Lemma 6.3.2 is not an equality in general. We will later see that this is often the case when X is a disc, $S = \mathcal{M}(k)$ and $S' = \mathcal{M}(l)$ for a finite wildly ramified extension l/k . One can also provide a characteristic-free example. Assume that $k = k^a$, $S = \mathcal{M}(k)$ and $X = S'$ is the unit disc over k , and let $x = s'$ be a point. It is easy to see that the fiber over (x, s') in X' contains a type 1 point x' . So, $\| \cdot \|_{x'} = 0$ by Sect. 6.2.3. On the other hand, if $r(x) > 0$ then $\| \cdot \|_x \neq 0$.

(ii) At first glance, the above examples are surprising because in the context of seminormed rings, Kähler seminorms are compatible with base changes by Lemma 4.2.6. However, we use structure sheafs provided with the spectral seminorms in the definition of Kähler seminorms, while the tensor seminorms do not have to be spectral. For example, if E/k and F/k are finite extensions of analytic fields such that $K = E \otimes_k F$ is a field, the tensor norm on K can be strictly larger than the valuation. As we will prove below, this phenomenon may only happen when the extensions are wild.

6.3.4 Universally Spectral Norms

Remark 6.3.3(ii) motivates the following definition. We say that a Banach k -algebra \mathcal{A} is *spectral* if its norm is power-multiplicative. In other words, $| \cdot |_{\mathcal{A}}$ coincides with the spectral seminorm. We say that a Banach k -algebra is *universally spectral* if $\mathcal{A} \widehat{\otimes}_k l$ provided with the tensor seminorm is spectral for any extension of analytic fields l/k .

Remark 6.3.5. (i) The condition that a Banach algebra is spectral is an analogue of reducedness in the usual ring theory. Thus, universal spectrality can be viewed as an analogue of geometric reducedness over a field.

(ii) For comparison, we note that a point $x \in X$ was called universal by Poineau, see [32, Definition 3.2], if the tensor norm on $\mathcal{H}(x) \widehat{\otimes}_k l$ is multiplicative for any l/k (originally, universal norms were called peaked, see [4, Sect. 5.2]). The algebraic analogue of the property that a Banach field is universal over k is geometrical integrality.

6.3.6 Defectless Case

One can show that an algebraic extension is universally spectral if and only if it is almost tame, but this will be worked out elsewhere. Here we only check this for defectless extensions.

Lemma 6.3.7. *Assume that K/k is a finite defectless extension of analytic fields. Then K is universally spectral over k if and only if K/k is tame.*

Proof. Throughout the proof, given an analytic k -field l we set $L = K \otimes_k l$ and provide it with the tensor seminorm $\| \cdot \|_L$. Let $\tilde{k}_{\text{gr}} = \bigoplus_{r>0} \tilde{k}_r$ be the $\mathbf{R}_{>0}$ -graded reduction of k , and define $\tilde{L}_{\text{gr}}, \tilde{K}_{\text{gr}}$ and \tilde{l}_{gr} similarly. Also, let \tilde{L}'_{gr} be the $\mathbf{R}_{>0}$ -graded ring associated with the filtration on L induced by $\| \cdot \|_L$; in particular, $\tilde{L}'_{\text{gr}} = \tilde{L}_{\text{gr}}$ only when $\| \cdot \|_L$ coincides with the spectral seminorm. Note that $\| \cdot \|_L$ is not spectral if and only if $\|x^n\|_L < \|x\|_L^n$ for some x and large n , and this happens if and only if \tilde{x} is a homogeneous nilpotent element of \tilde{L}'_{gr} . Thus, $\| \cdot \|_L$ is spectral if and only if the graded ring \tilde{L}'_{gr} is reduced in the sense of graded commutative algebra of [37, Sect. 1]. [This also reinforces Remark 6.3.5(i).]

Since K/k is defectless, K possesses an orthogonal k -basis a_1, \dots, a_n . Then, this basis is also an orthogonal basis of L over l and hence $\tilde{a}_1, \dots, \tilde{a}_n$ is a basis of both \tilde{K}_{gr} over \tilde{k}_{gr} and \tilde{L}'_{gr} over \tilde{l}_{gr} . In particular, we obtain that $\tilde{L}'_{\text{gr}} = \tilde{K}_{\text{gr}} \otimes_{\tilde{k}_{\text{gr}}} \tilde{l}_{\text{gr}}$.

Now, let us prove the lemma. The extension K/k is tame if and only if \tilde{K}/\tilde{k} is separable and $|K^\times|/|k^\times|$ has no p -torsion. By [14, Proposition 2.10] the latter happens if and only if the extension of graded fields $\tilde{K}_{\text{gr}}/\tilde{k}_{\text{gr}}$ is separable (in the graded setting). This implies that for any extension of analytic fields l/k , the graded \tilde{l}_{gr} -algebra $\tilde{K}_{\text{gr}} \otimes_{\tilde{k}_{\text{gr}}} \tilde{l}_{\text{gr}}$ is reduced, and as we saw above this happens if and only if the tensor seminorm on $K \otimes_k l$ is spectral.

Conversely, if K/k is wild, then we showed in the above paragraph that $\tilde{K}_{\text{gr}}/\tilde{k}_{\text{gr}}$ is not separable. For simplicity take $l = \widehat{k^a}$ (in fact, $l = K$ would suffice). Then \tilde{l}_{gr} is an algebraically closed graded field and $\tilde{K}_{\text{gr}} \otimes_{\tilde{k}_{\text{gr}}} \tilde{l}_{\text{gr}}$ is an inseparable finite \tilde{l}_{gr} -algebra. By [14, 1.14.3], this implies that $\tilde{K}_{\text{gr}} \otimes_{\tilde{k}_{\text{gr}}} \tilde{l}_{\text{gr}}$ is not reduced, and hence $\| \cdot \|_L$ is not spectral. □

Corollary 6.3.8. *Any tame extension is universally spectral.*

Proof. This follows from the lemma since any tame extension is defectless. □

6.3.9 Residually Tame Morphisms

We say that a morphism $g : S' \rightarrow S$ is *residually tame* (resp. *residually unramified*) at $s' \in S'$ if the extension of completed residue fields $\mathcal{H}(s')/\mathcal{H}(s)$, where $s = g(s')$, is finite and tame (resp. unramified). A morphism is *residually tame* or *residually unramified* if it is so at all points of the source.

6.3.10 Compatibility with Base Changes

In view of Remark 6.3.3(i), one has to restrict base change morphisms in order to ensure compatibility of Kähler seminorms with the base change. Here is a natural way to impose such a restriction.

Theorem 6.3.11. *Let $f : X \rightarrow S$ and $g : S' \rightarrow S$ be two morphisms of analytic k -spaces. Assume that for any point $s' \in S'$ with $s = g(s')$ the field $\mathcal{H}(s')$ is finite and universally spectral over $\mathcal{H}(s)$ (by Corollary 6.3.8 this includes the case of a residually tame g). Then the pullback of the Kähler seminorm on $\Omega_{X/S}$ coincides with the Kähler seminorm on $\Omega_{X \times_S S'/S'}$.*

Proof. Fix a point $x' \in X' = X \times_S S'$ and let $x \in X, s' \in S'$ and $s \in S$ be its images. Set also $K = \mathcal{H}(s), K' = \mathcal{H}(s'), L = \mathcal{H}(x)$, and $L'_1 = \mathcal{H}(x')$. We should prove that if $\phi \in \Omega_{X/S,x}$ and $\phi' \in \Omega_{X'/S',x'}$ is its pullback then $\|\phi\|_x = \|\phi'\|_{x'}$. By Corollary 6.1.9, the values of the seminorms can be computed at $\widehat{\Omega}_{L/K}$ and $\widehat{\Omega}_{L'_1/K'}$, respectively. This reduces the question to proving that the map $\psi : \widehat{\Omega}_{L/K} \otimes_L L'_1 \rightarrow \widehat{\Omega}_{L'_1/K'}$ is an isometry.

Consider the seminormed ring $L' = L \otimes_K K'$ and note that $\lambda : \widehat{\Omega}_{L/K} \otimes_L L' \xrightarrow{\sim} \widehat{\Omega}_{L'/K'}$ is an isometry by Lemma 4.2.6. In addition, L' is reduced since it is spectral, and L' is a finite L -algebra since K'/K is finite. Hence $L' \xrightarrow{\sim} \prod_{i=1}^n L'_i$, where L'_i are extensions of L and L'_1 is as defined earlier. We claim that this isomorphism is also an isometry. Indeed, the right-hand side is provided with the sup seminorm, hence this follows from [8, Theorem 3.8.3/7]. Thus, $\Omega_{L'/K'} = \prod_{i=1}^n \Omega_{L'_i/K'}$ and hence $\psi' : \widehat{\Omega}_{L'/K'} \otimes_{L'} L'_1 \rightarrow \widehat{\Omega}_{L'_1/K'}$ is an isometric isomorphism. Therefore, the composition $\psi = \psi' \circ (\lambda \otimes_{L'} L'_1)$ is an isometry. □

Let us record the most important particular cases of the theorem.

Corollary 6.3.12. *Assume that $X \rightarrow S$ is a morphism of k -analytic spaces, $s \in S$ is a point and X_s is the fiber over s . Then $\|\cdot\|_{\Omega_{X_s/S}}$ equals to the restriction of $\|\cdot\|_{\Omega_{X/S}}$ onto the fiber.*

Corollary 6.3.13. *Assume that $f : X \rightarrow S$ is a morphism of k -analytic spaces and l/k is a finite extension such that l is universally spectral over k (e.g., l/k is tame). Set $X_l = X \otimes l$ and $S_l = Y \otimes l$. Then $\|\cdot\|_{\Omega_{X_l/S_l}}$ is the pullback of $\|\cdot\|_{\Omega_{X/S}}$.*

Remark 6.3.14. (i) In fact, instead of finiteness of l/k it suffices to assume that the extension is algebraic. The proof is based on passing to a colimit and completing and we leave the details to the interested reader.

(ii) In particular, the Kähler seminorm is preserved when we replace k with its completed tame closure. In principle, this reduces the study of Kähler seminorms to the case of a tamely closed ground field k .

6.3.15 Geometric Kähler Seminorm

Given a morphism $f : X \rightarrow S$, let $\bar{f} : \bar{X} \rightarrow \bar{S}$ denote its ground field extension with respect to \bar{k}^a/k and let g denote the morphism $\bar{X} \rightarrow X$. Provide $g_* \Omega_{\bar{X}/\bar{S}}$ with the pushout quasi-norm $g_*(\|\cdot\|_{\Omega_{\bar{X}/\bar{S}}})$ (see Sect. 3.1.9) and let $\|\cdot\|_{\bar{\Omega}_{X/S}}$ be the quasi-norm it induces on $\Omega_{X/S}$ via the embedding $\Omega_{X/S} \hookrightarrow g_* \Omega_{\bar{X}/\bar{S}}$. By definition, if $\phi \in \Omega_{X/S}(U)$, then $\|\phi\|_{\bar{\Omega},V} = \|g^* \phi\|_{\Omega, g^{-1}(V)}$. It follows easily that $\|\cdot\|_{\bar{\Omega},X/S}$ is an

analytic seminorm and $\|\phi\|_{\overline{\Omega}_x} = \|g^*\phi\|_{\Omega_{x'}}$ for any $x' \in \overline{X}$ and $x = g(x')$. We call $\|\cdot\|_{\overline{\Omega}_{X/S}}$ the *geometric Kähler seminorm* on $\Omega_{X/S}$ and use $\overline{\Omega}$ in the notation to stress that it is geometric.

Absolutely in the same way one defines the *geometric Kähler seminorm* $\|\cdot\|_{(\overline{\Omega}_{X/S})^{\otimes m}}$ on the sheaf of pluriforms $(\Omega_{\overline{X/S}}^l)^{\otimes m}$.

Remark 6.3.16. If k possesses non-trivial wild extensions, Kähler seminorms on k -analytic spaces can behave rather weird (see Sects. 6.2.4 and 6.2.5). So, in this case it is often more useful to work with the geometric Kähler seminorm.

7 PL Subspaces

Main results about Kähler seminorms are related to the PL structure of analytic spaces, so in the current section we describe what this structure is.

7.1 Invariants of a Point

We start with recalling basic results about invariants t and s associated with points of k -analytic spaces, see [4, Sect. 9]. To stress the valuative-theoretic origin of these invariants we prefer to denote them F and E , the transcendental analogues of e and f .

7.1.1 Invariants of Extensions of Valued Fields

To any extension of valued fields L/K one can associate two cardinals that measure the “transcendence size” of the extension: the *residual transcendence degree* $F_{L/K} = \text{tr.deg.}_{\widetilde{K}}(\widetilde{L})$ and the *rational rank* $E_{L/K} = \dim_{\mathbf{Q}}(|L^\times|/|K^\times| \otimes_{\mathbf{Z}} \mathbf{Q})$. Both invariants are additive in towers of extensions $L'/L/K$, i.e. $E_{L'/K} = E_{L'/L} + E_{L/K}$ and $F_{L'/K} = F_{L'/L} + F_{L/K}$.

7.1.2 Invariants of Points

For a point x of a k -analytic space X we define the *residual transcendence degree* $F_{X,x} = F_{\mathcal{H}(x)/k}$ and the *rational rank* $E_{X,x} = E_{\mathcal{H}(x)/k}$. Usually, we will omit the space X in this notation. Additivity of the invariants can now be expressed as follows. Assume that $f : X \rightarrow Y$ is a morphism of k -analytic spaces, $x \in X$ is a point with $y = f(x)$, and the $\mathcal{H}(y)$ -analytic space $Z = X \times_Y \mathcal{M}(\mathcal{H}(y))$ is the fiber over y . Then $F_{X,x} = F_{Y,y} + F_{Z,x}$ and $E_{X,x} = E_{Y,y} + E_{Z,x}$. In particular, if f is finite, then $F_{X,x} = F_{Y,y}$ and $E_{X,x} = E_{Y,y}$.

7.1.3 Classification of Points on a Curve

Starting with [4, Sect. 1.4.4] and [5, Sect. 3.6], points of k -analytic curves are classified into four types: (1) $\mathcal{H}(x) \subseteq \widehat{k}^a$, (2) $F_x = 1$, (3) $E_x = 1$, (4) the rest. In all cases, it is easy to see that $E_x + F_x \leq 1$, i.e., $E_x = 0$ for type 2 points, $F_x = 0$ for type 3 points, and $E_x = F_x = 0$ for type 1 and 4 points.

If x is of type 2 or 3, then the following three claims hold: $\widetilde{\mathcal{H}}(x)$ is finitely generated over \widehat{k} , $|\mathcal{H}(x)^\times|$ is finitely generated over $|k^\times|$, if k is stable then $\mathcal{H}(x)$ is stable. The first two are simple, and we refer to [38, Corollary 6.3.6] for the stability theorem. Note that all three claims can fail for types 1 and 4.

7.1.4 Monomial Points

Fibering X by curves and using the additivity of E and F and induction on dimension, it is easy to see that any point $x \in X$ satisfies the inequality $E_x + F_x \leq \dim_x(X)$. A point $x \in X$ is called *monomial* or *Abhyankar* if $F_x + E_x = \dim_x(X)$. The set of all monomial points of X will be denoted X^{mon} .

- Remark 7.1.5.* (i) Monomial points are adequately controlled by the invariants E_x and F_x . This often makes the work with them much easier than with general points. For example, see Corollary 7.2.5 below.
- (ii) Analogues of monomial points in the theory of Riemann–Zariski spaces are often called Abhyankar valuations. They are much easier to work with too; for example, local uniformization is known for such points.

7.2 PL Subspaces

The set X^{mon} is huge and, at first glance, may look a total mess when $\dim(X) > 1$. Nevertheless, it possesses a natural structure of an ind-PL space that we are going to recall.

7.2.1 The Model Case

Recall that points of the affine space

$$\mathbf{A}_k^n = \bigcup_r \mathcal{M}(k\{r^{-1}t_1, \dots, r^{-1}t_n\})$$

with coordinates t_1, \dots, t_n can be identified with real semivaluations on $k[t_1, \dots, t_n]$ that extend the valuation of k . The semivaluations that do not vanish at t_1, \dots, t_n form the open subspace \mathbf{G}_m^n . Thus, for each tuple $r \in \mathbf{R}_{>0}^n$ the generalized gauss

valuation $|\cdot|_{\underline{t}}$ (see Sect. 5.7.3) defines a point $p_{\underline{t}} \in \mathbf{G}_m^n$, and the correspondence $\underline{t} \mapsto p_{\underline{t}}$ provides a topological embedding $\alpha : \mathbf{R}_{>0}^n \hookrightarrow \mathbf{G}_m^n$. Let S be the image of α .

- Remark 7.2.2.* (i) One often calls S the *skeleton* of \mathbf{G}_m^n ; this terminology is justified by (ii) and (iii) below. Sometimes one refers to the points of S as *\underline{t} -monomial valuations* because they are determined by their restriction to the monoid $\prod_{i=1}^n t_i^{\mathbf{N}}$.
- (ii) Any semivaluation $x \in \mathbf{G}_m^n$ is dominated by the \underline{t} -monomial valuation $p_{|\underline{t}(x)|}$. The map $x \mapsto p_{|\underline{t}(x)|}$ is a retraction $r : \mathbf{G}_m^n \twoheadrightarrow S$. Moreover, Berkovich constructs in [4, Sect. 6] a deformational retraction of \mathbf{G}_m^n onto S whose level at 1 is the above map.
- (iii) All \underline{t} -monomial valuations are distinguished by invertible functions (in fact, by monomials t^a), hence they are incomparable with respect to the domination. Thus, S is the set of all points of \mathbf{G}_m^n that are maximal with respect to the domination. In particular, it is independent of the coordinates.

7.2.3 Monomial Charts

By a monomial chart of an analytic space X we mean a morphism $f : U \rightarrow \mathbf{G}_m^n$ such that U is an analytic domain in X and f has zero-dimensional fibers. Clearly, f is determined by invertible functions $t_1, \dots, t_n \in \mathcal{O}_X^\times(U)$. A point $x \in U$ of the chart is called *\underline{t} -monomial* or *f -monomial* if the restriction of $|\cdot|_x$ to $k[t_1, \dots, t_n]$ is monomial. In this case, we also say that t_1, \dots, t_n is a *family of monomial parameters* at x . Note that $\mathcal{H}(f(x)) = \widehat{k(\underline{t})}$ is provided with a generalized Gauss valuation and $\mathcal{H}(x)$ is its finite extension.

In addition, we say that the chart is residually tame or unramified (at a point $x \in U$) if the morphism f is so, see Sect. 6.3.9. In this case we also say that the family of monomial parameters t_1, \dots, t_n is *residually tame* or *unramified* (at x). The following result shows that monomial charts adequately describe the whole X^{mon} .

Lemma 7.2.4. *A point $x \in X$ is monomial if and only if there exists a monomial chart $f : U \rightarrow \mathbf{G}_m^n$ such that x is f -monomial.*

Proof. Set $L = \mathcal{H}(x)$. If there exists a monomial chart f such that $f(x)$ is a monomial point, then the induced map $f : U \rightarrow \mathbf{G}_m^n$ has zero-dimensional fibers and hence L is finite over $K = \mathcal{H}(y)$, where $y = f(x)$. Since $E_y + F_y = n$, we obtain that $E_x + F_x = n$, i.e. x is monomial.

Conversely, assume that the sum of $E = E_x$ and $F = F_x$ equals to $n = \dim_x(X)$. Choose $t_1, \dots, t_F \in \kappa_G(x)$ such that $\tilde{t}_1, \dots, \tilde{t}_F$ is a transcendence basis of $\widetilde{L}/\widetilde{k}$, and choose $\{t_{F+1}, \dots, t_n\}$ such that its image in $(|L^\times|/|k^\times|) \otimes \mathbf{Q}$ is a basis. Then the valuation on $K = \widehat{k(t_1, \dots, t_n)}$ is a generalized Gauss valuation. Take an analytic domain $U \subseteq X$ containing x such that $\dim(U) = n$ and t induces a morphism $f : U \rightarrow \mathbf{G}_m^n$. Then $y = f(x)$ is a monomial point of \mathbf{G}_m^n (even a point of its skeleton). Consider the fiber $U_y = f^{-1}(y)$. By [15, Corollary 8.4.3] $\dim_y(U_y) = 0$, hence using [12, Theorem 4.9] we can shrink U around y so that f has zero-dimensional fibers.

□

Corollary 7.2.5. *Assume that $x \in X$ is a monomial point, then*

- (1) *The ring $\mathcal{O}_{X_G, x}$ is Artin. In particular, if X is reduced, then $m_{G, x} = 0$ and $\mathcal{O}_{X_G, x} = \kappa_G(x)$.*
- (2) *The extension $\widetilde{\mathcal{H}(x)}/\widetilde{k}$ is finitely generated.*
- (3) *The group $|\mathcal{H}(x)^\times|/|k^\times|$ is finitely generated.*
- (4) *If k is stable, then $\mathcal{H}(x)$ is stable.*

Proof. By Lemma 7.2.4 there exists a chart $f : U \rightarrow \mathbf{G}_m^n$ that takes x to a point y corresponding to a generalized Gauss valuation. Note that $\mathcal{O}_{X_G, x}$ is finite over $\mathcal{O}_{Y_G, y}$. All four properties are satisfied by y , hence they also hold for x . □

We can now strengthen Lemma 7.2.4 in the case of an algebraically closed k .

Corollary 7.2.6. *Assume X is a k -analytic space, k algebraically closed and $x \in X$ is a monomial point. Then there exists a monomial chart f such that x is f -monomial and f is residually unramified at x .*

Proof. The proof repeats that of Lemma 7.2.4, but we will choose t_i more carefully. Set $L = \mathcal{H}(x)$, $E = E_{L/k}$, $F = F_{L/k}$ and $n = E + F$. Since $k = k^a$ the field \widetilde{k} is algebraically closed and the group $|k^\times|$ is divisible. In particular, the extension $\widetilde{L}/\widetilde{k}$ is separable and the group $|L^\times|/|k^\times|$ is torsion free. Since $\widetilde{L}/\widetilde{k}$ is finitely generated by Corollary 7.2.5(2), it possesses a separable transcendence basis. Choose $t_1, \dots, t_F \in \kappa_G(x)$ so that $\widetilde{t}_1, \dots, \widetilde{t}_F$ is such a basis. Since the group $|L^\times|/|k^\times|$ is finitely generated by Corollary 7.2.5(3), it is isomorphic to \mathbf{Z}^E . Choose $t_{F+1}, \dots, t_n \in \kappa_G(x)$ so that their images form a basis of $|L^\times|/|k^\times|$.

The elements t_1, \dots, t_n define a morphism $f : U \rightarrow \mathbf{G}_m^n$ for a small enough analytic domain U containing x , and we claim that f is as required. Indeed, it suffices to check that L is unramified over $K = \mathcal{H}(f(x)) = \overline{k(t_1, \dots, t_n)}$. By our choice $|L^\times| = |K^\times|$ and $\widetilde{L}/\widetilde{K}$ is separable. Since K is stable by Corollary 7.2.5(4), L/K is unramified. □

7.2.7 Skeletons of Monomial Charts

The set of all f -monomial points of U will be denoted $S(f)$ and called the *skeleton* of the chart. Note that $S(f)$ is nothing else but the preimage of the skeleton of \mathbf{G}_m^n under f . By Lemma 7.2.4, $X^{\text{mon}} = \cup_f S(f)$, where f runs over all monomial charts of X . We warn the reader that, in general, $S(f)$ does not have to be a retract of U .

7.2.8 R_S -PL Structures

We will usually abbreviate “piecewise linear” as PL. We refer to [7, Sect. 1] for the definition of an R_S -PL space Q for a ring $R \subseteq \mathbf{R}$ and its exponential module $S \subseteq \mathbf{R}_{>0}$. Here we only recall that Q has an atlas $\{P_i\}$ of R_S -PL polytopes, i.e. polytopes in $\mathbf{R}_{>0}^n$ given by finitely many inequalities $st_1^{e_1} \dots t_n^{e_n} \leq 1$ with $s \in S$, $e_i \in R$ and provided with the family of R_S -PL functions.

7.2.9 Rational PL-Subspaces

Absolute values of the coordinates of \mathbf{G}_m^n induce coordinates $t_i : S \rightarrow \mathbf{R}_{>0}$ on its skeleton S . The latter are unique up to the action of $\mathrm{GL}(n, \mathbf{Z}) \ltimes |k^\times|^n$ combined from the action of $\mathrm{GL}(n, \mathbf{Z})$ on $\prod_{i=1}^n t_i^{\mathbf{Z}}$ and rescaling the coordinates by elements of $|k^\times|$. Therefore, S acquires a canonical $\mathbf{Z}_{|k^\times|}$ -PL structure.

The following facts were proved by Ducros: (1) for any monomial chart $f : U \rightarrow \mathbf{G}_m^n$ the skeleton $S(f)$ possesses a unique \mathbf{Q}_H -PL structure such that the map $S(f) \rightarrow S$ is \mathbf{Q}_H -PL, see [11, Theorem 3.1], (2) for any two monomial charts f and g the intersection $S(f) \cap S(g)$ is \mathbf{Q}_H -PL in both $S(f)$ and $S(g)$, see [13, Theorem 5.1], and so $S(f) \cup S(g)$ acquires a natural \mathbf{Q}_H -PL structure and the whole X^{mon} acquires a natural structure of an ind- \mathbf{Q}_H -PL space. By a *rational \mathbf{Q}_H -PL subspace* of X we mean a subset of the form $\cup_{i=1}^n S(f_i)$ with its induced \mathbf{Q}_H -PL structure.

7.2.10 Integral Structure

One may wonder if the \mathbf{Q}_H -PL structure on a rational PL subspace P of X can be refined to an integral one. Ducros and Thuillier showed in [16] that this is indeed the case: there is a \mathbf{Z}_H -PL structure on P such that a function $f : P \rightarrow \mathbf{R}_{>0}$ is \mathbf{Z}_H -PL if and only if G -locally it is of the form $r|f|$, where f is an analytic function on X and $r \in H$ (see [16, 3.7] and note that $H^{\mathbf{Q}} = H$). Obviously, such a structure is unique but consistency of the definition requires an argument. In addition, they show that the associated \mathbf{Q}_H -PL space, obtained by adjoining integral roots of all \mathbf{Z}_H -PL functions on P , coincides with the original \mathbf{Q}_H -PL space. In the sequel, we provide P with this \mathbf{Z}_H -PL structure and call it a *\mathbf{Z}_H -PL subspace* of X .

7.3 Semistable Formal Models

7.3.1 Semistable Formal Schemes

We say that a formal k° -scheme is *strictly semistable* if locally it admits an étale morphism to a formal scheme of the form $\mathfrak{Z}_{n,a} = \mathrm{Spf}(k^\circ\{T_0, \dots, T_n\}/(T_0 \dots T_n - a))$ with $0 \neq a \in k^\circ$. A formal k° -scheme is called *semistable* if it is étale-locally strictly semistable.

Remark 7.3.2. (i) Sometimes one does not require that $a \neq 0$, thereby obtaining a wider class of semistable formal schemes. If \mathfrak{X} is semistable in this sense, then \mathfrak{X}_η is quasi-smooth if and only if one can find charts with $a \neq 0$. Since we will only be interested in formal models of quasi-smooth spaces, we use the definition that includes the condition $a \neq 0$.

- (ii) If the valuation is discrete and a is a uniformizer, one often considers schemes $\mathfrak{Z} = \mathrm{Spf}(k^\circ\{T_0, \dots, T_n\}/(T_0^{l_0} \dots T_n^{l_n} - a))$. They are regular with SNC closed fiber. If $l_i > 1$, then the closed fiber is not reduced and hence \mathfrak{Z} is not semistable. Note also that if $(l_1, \dots, l_n) \in \tilde{k}^\times$ then \mathfrak{Z} is log smooth, see Remark 7.3.4 below.

7.3.3 Skeletons Associated to Semistable Formal Models

To any semistable formal k° -scheme \mathfrak{X} with generic fiber $X = \mathfrak{X}_\eta$ Berkovich associated in [6, Sect. 5] the skeleton $S(\mathfrak{X}) \subset X$ and constructed a deformation retraction $X \rightarrow S(\mathfrak{X})$. Moreover, these constructions are compatible with any étale morphism $\phi : \mathfrak{Y} \rightarrow \mathfrak{X}$, i.e. $S(\mathfrak{Y}) = \phi_\eta^{-1}(S(\mathfrak{X}))$ and the retraction is compatible with ϕ_η , see [6, Theorem 5.2(vii)].

Since any semistable formal k° -scheme is connected with semistable formal schemes $\mathfrak{Z}_{n,a}$ by a zigzag of two étale morphisms, description of the skeleton and the retraction reduces to the model case, and the latter is induced from \mathbf{G}_m^n . Namely, let \mathbf{G}_m^n be the n -dimensional torus with coordinates T_1, \dots, T_n . The affinoid subdomain given by $|T_1 \dots T_n| \geq |a|$ and $|T_i| \leq 1$ for $1 \leq i \leq n$ equals to $Z = \mathcal{M}(k\{T_0, \dots, T_n\}/(T_0 \dots T_n - l))$ and hence can be identified with the generic fiber of $\mathfrak{Z} = \mathfrak{Z}_{n,a}$. Then T_1, \dots, T_n give rise to the monomial chart $f : Z \hookrightarrow \mathbf{G}_m^n$ with skeleton $S(f) = S(\mathbf{G}_m^n) \cap Z$, and $S(\mathfrak{Z})$ coincides with $S(f)$. The retraction of Z onto $S(\mathfrak{Z})$ is the restriction of the retraction of \mathbf{G}_m^n onto $S(\mathbf{G}_m^n)$.

Remark 7.3.4. (i) Slightly more generally, Berkovich makes the above constructions for *polystable* models, i.e. those models that are étale-locally isomorphic to products of semistable ones. In fact, this can be extended further to log smooth formal models with trivial generic log structure—these are formal schemes that étale-locally admit smooth morphisms to formal schemes of the form $\mathrm{Spf}(k^\circ\{P/\Pi\})$, where Π and P are sharp fs monoids, $\alpha : \Pi \hookrightarrow k^\circ \setminus \{0\}$ is an embedding such that the composition $\Pi \rightarrow |k^\circ \setminus \{0\}|$ is injective, $\phi : \Pi \rightarrow P$ is an injective homomorphism with no p -torsion in the cokernel and such that $\Pi^{\mathrm{gp}}P = P^{\mathrm{gp}}$, and $k^\circ\{P/\Pi\}$ is the quotient of $k^\circ\{P\}$ by the ideal generated by elements $\alpha(\pi) - \phi(\pi)$ for $\pi \in \Pi$. The details will be worked out elsewhere.

- (ii) The main motivation for considering log smooth formal models is that it is believed (at least by the author) that any quasi-smooth compact strictly analytic space possesses a log smooth formal model. This is absolutely open when $\mathrm{char}(\tilde{k}) > 0$, but one may hope to prove this by current techniques when $\mathrm{char}(\tilde{k}) = 0$. Currently, the latter is only known for a discretely valued k : desingularization of excellent formal schemes implies that there even exists a semistable formal model. If the \mathbf{Q} -rank of $|k^\times|$ is larger than one, then it is easy to give examples when semistable models do not exist, e.g. $\mathrm{Spf}(k\{at^{-1}, t\}) \times \mathrm{Spf}(k\{bt^{-1}, t\})$, where $|b| \notin |a|^\mathbf{Q}$ (see also [41, Remark 1.1.1(ii)]). The situation with polystable models is unclear: it is still an open combinatorial question

(in dimension at least 4) whether any log smooth formal model possesses a polystable refinement. A tightly related question was raised by Abramovich and Karu in [1, Sect. 8]: log smooth (resp. polystable) models are analogues of weakly semistable (resp. semistable) morphisms in the sense of [1, Sect. 0].

8 Metrization of Pluricanonical Forms

Throughout Sect. 8, X is assumed to be quasi-smooth of pure dimension n . In particular, $\Omega_X = \Omega_{X/k}^1$ is a locally free sheaf of rank n and the pluricanonical sheaves $\omega_X^{\otimes m} = (\bigwedge^n \Omega_X)^{\otimes m}$ are invertible.

8.1 Monomiality of Kähler Seminorms

8.1.1 Stalks at Monomial Points

Recall that if $k = k^a$ then any monomial point possesses a family of residually tame monomial parameters by Corollary 7.2.6. This allows to describe $\|\cdot\|_\omega$ as follows.

Theorem 8.1.2. *Assume that k is algebraically closed. Let X be a quasi-smooth k -analytic space of dimension n with a point $x \in X$, and let t_1, \dots, t_n be invertible elements of $\mathcal{O}_{X_G, x}$. Then $\|\frac{dt_1}{t_1} \wedge \dots \wedge \frac{dt_n}{t_n}\|_{\omega, x} \leq 1$ and the following conditions are equivalent:*

- (i) $\|\frac{dt_1}{t_1} \wedge \dots \wedge \frac{dt_n}{t_n}\|_{\omega, x} = 1$.
- (ii) $\frac{dt_1}{t_1}, \dots, \frac{dt_n}{t_n}$ form an orthonormal basis of $\widehat{\Omega}_{\mathcal{H}(x)/k}$.
- (iii) t is a family of residually tame monomial parameters at x .

Proof. Since $\|dt_i\|_{\Omega, x} \leq |t_i|_x$, it follows that $\|\frac{dt_1}{t_1} \wedge \dots \wedge \frac{dt_n}{t_n}\|_{\omega, x} \leq 1$. The equivalence (i) \iff (ii) is proved precisely as the equivalence (iii) \iff (iv) in the proof of Theorem 5.7.8. We will complete the proof by showing that (ii) \iff (iii).

Set $l = k(t)$ and $L = \hat{l}$. First, assume that t_i form a family of residually tame monomial parameters. Then $\frac{dt_1}{t_1}, \dots, \frac{dt_n}{t_n}$ is an orthonormal basis of $\widehat{\Omega}_{l/k}$ by Theorem 5.7.8. It remains to use that $\widehat{\Omega}_{l/k} = \widehat{\Omega}_{L/k}$ by Corollary 5.6.7, and $\widehat{\Omega}_{L/k} = \widehat{\Omega}_{\mathcal{H}(x)/k}$ because $\mathcal{H}(x)/L$ is tame.

Conversely, assume that $\frac{dt_1}{t_1}, \dots, \frac{dt_n}{t_n}$ is an orthonormal basis of $\widehat{\Omega}_{\mathcal{H}(x)/k}$. Since the isomorphism $\psi_{\mathcal{H}(x)/L/k} : \widehat{\Omega}_{L/k} \otimes_L \mathcal{H}(x) \rightarrow \widehat{\Omega}_{\mathcal{H}(x)/k}$ is non-expansive and $\|\frac{dt_i}{t_i}\|_{\Omega, L/k} \leq 1$, we obtain that $\psi_{\mathcal{H}(x)/L/k}$ is an isometry and $\frac{dt_1}{t_1}, \dots, \frac{dt_n}{t_n}$ is an orthonormal basis of $\widehat{\Omega}_{L/k}$. The module $\Omega_{\mathcal{H}(x)^\circ/k^\circ}^{\log}$ is almost torsion free by Theorem 5.5.18, hence $\mathcal{H}(x)/L$ is almost tame by Theorem 5.6.4.

It remains to show that t_1, \dots, t_n is a family of monomial parameters because then L is stable by Corollary 7.2.5(4) and hence $\mathcal{H}(x)/L$ is tame by Theorem 5.5.11(ii).

Note that t_1, \dots, t_n are algebraically independent over k because dt_1, \dots, dt_n are linearly independent in $\widehat{\Omega}_{\mathcal{H}(x)/k}$, hence t_1, \dots, t_n is a separable transcendence basis of l/k . Since $\Omega_{l/k}$ is finite-dimensional, the completion $\Omega_{l/k} \rightarrow \widehat{\Omega}_{l/k}$ is surjective, and comparing the dimensions we see that it is an isomorphism. Thus, $\Omega_{l/k} \otimes_l L \rightarrow \widehat{\Omega}_{L/k}$ is an isometric isomorphism by Corollary 5.6.7. This implies that $\frac{dt_1}{t_1}, \dots, \frac{dt_n}{t_n}$ is an orthonormal basis of $\Omega_{l/k}$, and hence the valuation on l is t -monomial by Theorem 5.7.8. \square

Corollary 8.1.3. *Assume that k is algebraically closed, X is quasi-smooth, elements $t_1, \dots, t_n \in \mathcal{O}_{X_{G,x}}$ form a family of residually tame monomial parameters at a point x , and $\mathcal{F} = (\Omega_X^l)^{\otimes m}$. Then*

$$B = \left\{ \left(\frac{dt_{i_1}}{t_{i_1}} \wedge \dots \wedge \frac{dt_{i_l}}{t_{i_l}} \right) \otimes \dots \otimes \left(\frac{dt_{j_1}}{t_{j_1}} \wedge \dots \wedge \frac{dt_{j_l}}{t_{j_l}} \right) \right\}$$

is an orthonormal basis of \mathcal{F}_x . In particular, any pluriform $\phi \in \Gamma((\Omega_X^l)^{\otimes m})$ can be represented as $\phi = \sum_{e \in B} \phi_e e$ locally at x , and the following equality holds

$$\|\phi\|_{(\Omega_X^l)^{\otimes m, x}} = \max_{e \in B} |\phi_e|_x.$$

Proof. By Theorem 8.1.2, $\frac{dt_1}{t_1}, \dots, \frac{dt_n}{t_n}$ is an orthonormal basis of $\Omega_{X,x}$. This reduces the claim to simple multilinear algebra. \square

Corollary 8.1.4. *Keep the assumptions of Corollary 8.1.3 and assume that $\mathcal{F} = \omega_X^{\otimes m}$. Then $e = (\frac{dt_1}{t_1} \wedge \dots \wedge \frac{dt_n}{t_n})^{\otimes m}$ is a basis of \mathcal{F}_x , and if $\phi = fe$ is the representation of a pluricanonical form ϕ at x then $\|\phi\|_{\omega^{\otimes m, x}} = |f|_x$.*

8.1.5 Piecewise Monomiality

Now, we can prove that norms of pluriforms induce \mathbf{Z}_H -PL functions on \mathbf{Z}_H -PL subspaces of X (see Sect. 7.2.9). In particular, when restricted to \mathbf{Z}_H -PL subspaces of X , spectral seminorm and Kähler seminorms demonstrate similar behavior, although their global behavior is very different.

Theorem 8.1.6. *Assume that X is a reduced k -analytic space and $\phi \in \Gamma((\Omega_X^l)^{\otimes m})$ is a pluriform on X , and consider the function $\|\phi\| : X \rightarrow \mathbf{R}_{\geq 0}$ that sends x to the value $\|\phi\|_{(\Omega_X^l)^{\otimes m, x}}$ of the geometric Kähler seminorm. Then for any \mathbf{Z}_H -PL subspace $P \subset X$ the restriction of $\|\phi\|$ onto P is a \mathbf{Z}_H -PL function.*

Proof. Since X is reduced, any monomial point satisfies $m_x = 0$ and hence X is quasi-smooth at x . Therefore, replacing X by a neighborhood of P we can assume that X is quasi-smooth.

First, we consider the case when $k = k^a$. By Theorem 9.4.8 that will be proved in the end of the paper, we can cover S by finitely many skeletons of residually tame (even unramified) monomial charts, hence it suffices to consider the case when P

itself is the skeleton of a residually tame monomial chart $f : U \rightarrow \mathbf{G}_m^n$ given by $t_1, \dots, t_n \in \Gamma(\mathcal{O}_U^\times)$. By Corollary 8.1.3, locally at a point $x \in P$ we can represent $\|\phi\|$ as the maximum of \mathbf{Z}_H -PL functions $|\phi_i|$. Hence $\|\phi\|$ is \mathbf{Z}_H -PL too.

Assume, now, that k is arbitrary. Set $\overline{X} = X \otimes_k k^a$ and let $\overline{\phi} \in \Gamma((\Omega_{\overline{X}}^1)^{\otimes m})$ be the pullback of ϕ . The preimage $\overline{P} \subset \overline{X}$ of P is a \mathbf{Z}_H -PL subspace: just take the charts of P and extend the ground field to k^a . Also, it follows from the description with charts that the map $\overline{P} \rightarrow \overline{X}$ is \mathbf{Z}_H -PL. Since the pullback of $\|\phi\|_{(\overline{\Omega}_X^1)^{\otimes m, x}}$ to \overline{P} is $\|\overline{\phi}\|_{(\Omega_{\overline{X}}^1)^{\otimes m, x}}$ and the latter function is \mathbf{Z}_H -PL by the case of an algebraically closed ground field, $\|\phi\|_{(\overline{\Omega}_X^1)^{\otimes m, x}}$ is a \mathbf{Z}_H -PL function as well. \square

Remark 8.1.7. Most probably, the assumption that X is reduced can also be removed. Also, it seems probable that the theorem holds for the Kähler seminorm too.

8.2 Maximality Locus of Pluricanonical Forms

Given a pluricanonical form ϕ , let M_ϕ be the maximality locus of the Kähler seminorm of ϕ and let \overline{M}_ϕ be the maximality locus of the geometric Kähler seminorm of ϕ . Our next aim is to study \overline{M}_ϕ ; recall that it is closed by Theorem 6.1.13 and Lemma 3.3.10. The main results of this section, including Theorems 8.2.4 and 8.2.9, do not hold for the maximality locus M_ϕ , at least when the residue field is not perfect; the counterexamples being as in Sect. 6.2.5 and Remark 6.2.6.

8.2.1 The Torus Case

We start with studying the standard pluricanonical form on a torus.

Lemma 8.2.2. *Let X be the k -analytic torus $\mathbf{G}_m^{n, \text{an}}$ with coordinates t_1, \dots, t_n . Consider the pluricanonical form $\phi = (\frac{dt_1}{t_1} \wedge \dots \wedge \frac{dt_n}{t_n})^{\otimes m}$. Then $\|\phi\|_{\overline{\omega}^{\otimes m, x}} \leq 1$ for any $x \in X$ and the equality takes place if and only if x is a generalized Gauss point (see Sect. 7.2.1). In particular, the maximum locus of $\|\phi\|_{\overline{\omega}^{\otimes m}}$ is the skeleton $\mathbf{R}_{>0}^n$ of X .*

Proof. This immediately follows from Theorem 8.1.2. \square

8.2.3 The Semistable Case

Recall that the skeleton $S(\mathfrak{X}) \subset X$ associated with a strictly semistable formal model has a natural structure of a simplicial complex.

Theorem 8.2.4. *Assume that X is a quasi-smooth compact strictly k -analytic space, $\phi \in \Gamma(\omega_X^{\otimes m})$ is a pluricanonical form on X , \mathfrak{X} is a strictly semistable formal model of X , and $S(\mathfrak{X}) \subset X$ is the skeleton associated with \mathfrak{X} . Then the maximality locus \overline{M}_ϕ is a union of faces of $S(\mathfrak{X})$.*

Proof. Set $\| \cdot \| = \| \cdot \|_{\overline{\omega}^{\otimes m}}$ for shortness. First, let us prove that $\overline{M}_\phi \subseteq S(\mathfrak{X})$. Working locally on \mathfrak{X} we can assume that there exists an étale morphism

$$g : \mathfrak{X} \rightarrow \mathfrak{Y} = \text{Spf}(k^\circ \{t_0, \dots, t_n\} / (t_0 \dots t_n - a)),$$

where $0 \neq a \in k^\circ$. Since $Y = \mathfrak{Y}_\eta$ is a domain in \mathbf{G}_m^n with coordinates t_1, \dots, t_n (see Sect. 7.3.3), $e = (\frac{dt_1}{t_1} \wedge \dots \wedge \frac{dt_n}{t_n})^{\otimes m}$ is a nowhere vanishing pluricanonical form on Y and therefore $\phi = he$ for a function $h \in \Gamma(\mathcal{O}_X)$. By [6, Theorem 5.3.2(ii)], the retraction $r : X \rightarrow S(\mathfrak{X})$ is compatible with domination, in particular, $|h|_x \leq |h|_{r(x)}$. Since $\|\phi\|_x = |h|_x \|e\|_x$, it remains to prove that $\|e\|_x \leq \|e\|_{r(x)}$ and the equality holds if and only if $x = r(x)$, i.e. $x \in S(\mathfrak{X})$.

When working with e , we can also view it as a form on Y . Let $g : X \rightarrow Y$ be the generic fiber of g . Since g is étale, if $x \in X$ and $y = g(x)$ then the extension $\mathcal{H}(x)/\mathcal{H}(y)$ is unramified by [6, Lemma 1.6]. So, Theorem 6.3.11 and Lemma 6.3.7 imply that $\|e\|_x = \|e\|_y$. By Corollary 8.2.2, $\|e\|_y \leq 1$ for any point $y \in Y$ and the equality holds precisely for the points of $S(\mathfrak{Y})$. So, $\|e\|_x \leq 1$ for any $x \in X$ and the equality holds precisely for the points of $g^{-1}(S(\mathfrak{Y})) = S(\mathfrak{X})$.

It remains to show that if Δ is a face of $S(\mathfrak{X})$ then either $\Delta \subset \overline{M}_\phi$ or the interior Δ° is disjoint from \overline{M}_ϕ . This claim is local on \mathfrak{X} so we can assume that $\mathfrak{X} = \text{Spf}(\mathcal{A}^\circ)$ is affine and there is an étale morphism $g : \mathfrak{X} \rightarrow \mathfrak{Y}$ as above. In particular, we can assume that $\phi = he$, as above, and so $\|\phi\|_x = |h|_x$ at any point $x \in S(\mathfrak{X})$. Note that \mathfrak{X} is pluri-nodal in the sense of [6, Sect. 1], hence $|h|_{\mathcal{A}} \in |k^\times|$ by [6, Proposition 1.4]. Thus, multiplying ϕ by an element of k^\times we can achieve that $|h|_{\mathcal{A}} = 1$, and then $h \in \mathcal{A}^\circ$ is a function on \mathfrak{X} . Note that Δ° is the preimage in $S(\mathfrak{X})$ of a point $\mathfrak{r} \in \mathfrak{X}$. If $\tilde{h}(\mathfrak{r}) = 0$, then $|h|_x < 1$ at any point in the preimage of \mathfrak{r} in X , and hence $\Delta^\circ \cap \overline{M}_\phi = \emptyset$. Otherwise, $|h|_x = 1$ for any x as above, hence $\Delta^\circ \subseteq \overline{M}_\phi$ and by the closedness of \overline{M}_ϕ we obtain that $\Delta \subseteq \overline{M}_\phi$. \square

Remark 8.2.5. (i) Theorem 8.2.4 is an analogue of [28, Theorem 4.5.5], though it applies to a wider context (e.g., X is only assumed to be quasi-smooth). Note, however, that these results consider different seminorms when $\text{char}(k) > 0$ (see Sect. 8.3 below).

(iii) The lemma can be extended to general semistable models at cost of considering generalized simplicial complexes. Moreover, it should extend to arbitrary log smooth formal models, once the foundations are set (see Remark 7.3.4).

8.2.6 Residually Tame Coverings

Recall that residual tameness was defined in Sect. 6.3.9.

Lemma 8.2.7. *Assume that X is a quasi-smooth compact strictly k -analytic space admitting a residually tame quasi-étale covering $Y \rightarrow X$ such that Y possesses a strictly semistable formal model. Then for any pluricanonical form $\phi \in \Gamma(\omega_X^{\otimes m})$ on X the geometric maximality locus \bar{M}_ϕ is a compact \mathbf{Z}_H -PL subspace of X .*

Proof. Let $\psi \in \Gamma(\omega_Y)$ be the pullback of ϕ . By Theorem 6.3.11, the norm functions $\|\phi\|_{\omega^{\otimes m}}$ and $\|\psi\|_{\omega^{\otimes m}}$ are compatible, and hence $\bar{M}_\phi = f(\bar{M}_\psi)$. It remains to recall that \bar{M}_ψ is a compact \mathbf{Z}_H -PL subspace by Theorem 8.2.4, and hence its image under f is a \mathbf{Z}_H -PL subspace by [16, Proposition 2.1]. □

8.2.8 Residue Characteristic Zero

In order to use the previous lemma, one should construct an appropriate covering $Y \rightarrow X$. The author conjectures that any quasi-smooth strictly analytic space possesses a quasi-net of analytic subdomains that admit a semistable model (this is an analogue of the local uniformization conjecture in the desingularization theory). However, this seems to be out of reach when $\text{char}(\tilde{k}) > 0$. Even the case of $\text{char}(\tilde{k}) = 0$, which should be relatively simple, is missing in the literature. We can avoid dealing with it here in view of the following generalization of a theorem of U. Hartl to ground fields with non-discrete valuations, see [41, Theorem 3.4.1]: there exist a finite extension l/k and a quasi-étale surjective morphism $Y \rightarrow X_l = X \otimes_k l$ such that Y possesses a strictly semistable formal model. The theorem does not provide any control on residual tameness of f , but it is automatic whenever $\text{char}(\tilde{k}) = 0$.

Theorem 8.2.9. *Assume that $\text{char}(\tilde{k}) = 0$ and X is a quasi-smooth compact strictly k -analytic space. Then for any pluricanonical form $\phi \in \Gamma(\omega_X^{\otimes m})$ on X , the maximality locus M_ϕ is a compact \mathbf{Z}_H -PL subspace of X .*

Proof. In this case, there is no difference between $\|\cdot\|_{\omega^{\otimes m}}$ and $\|\cdot\|_{\bar{\omega}^{\otimes m}}$. Take Y and l as in [41, Theorem 3.4.1], then the composition $Y \rightarrow X \otimes_k l \rightarrow X$ is a quasi-étale covering, which is automatically residually tame. It remains to use Lemma 8.2.7. □

We say that a point $x \in X$ is *divisorial* if x is monomial and $E_{\mathcal{H}(x)/k} = 0$. Thus, x is divisorial if and only if $\text{tr.deg.} \widetilde{\mathcal{H}(x)/\tilde{k}} = \dim_x(X)$. We will not need this, but it is easy to see that if X is strictly analytic then x is divisorial if and only if there exists a formal model \mathfrak{X} such that x is the preimage of a generic point of \mathfrak{X}_s under the reduction map.

Corollary 8.2.10. *Keep the assumptions of Theorem 8.2.9 and let X^{div} be the set of divisorial points of X . Then $\|\phi\|_{\omega^{\otimes m}} = \max_{x \in X^{\text{div}}} \|\phi\|_{\omega^{\otimes m}, x}$.*

Proof. By Theorem 8.2.9, the maximality locus M_ϕ is a compact \mathbf{Z}_H -PL subspace. By Theorem 8.1.6, the restriction of $\|\phi\|$ on M_ϕ is a \mathbf{Z}_H -PL function, hence it achieves maximum at a \mathbf{Z}_H -rational point x . Any such x is a divisorial point of X . □

8.3 Comparison with the Weight Norm of Mustařă–Nicaise

We conclude Sect. 8 by comparing $\|\cdot\|_\omega$ with the weight norm à la Mustařă and Nicaise, see [28]. Unless said to the contrary, k is assumed to be discretely valued.

8.3.1 Weight Seminorm

Assume that K/k is a separable finitely generated extension of real-valued fields of transcendence degree n such that $\text{tr.deg.}(\tilde{K}/\tilde{k}) = n$. Note that K is discretely valued and such extensions correspond to divisorial valuations. Let us recall how a norm on $\omega_{K/k} = \Omega_{K/k}^n$ is defined in [28]. We will call it the *weight norm* and denote $\|\cdot\|_{\text{wt}}$. Fix $t_1, \dots, t_n \in K^\circ$ such that $\tilde{t}_1, \dots, \tilde{t}_n$ is a transcendence basis of \tilde{K}/\tilde{k} . We claim that replacing t_i with elements t'_i such that $|t_i - t'_i| < 1$ one can in addition achieve that t_1, \dots, t_n is a separable transcendence basis of K . To prove this we will use the observation that the latter happens if and only if dt_1, \dots, dt_n is a basis of $\Omega_{K/k}$. Choose a separable transcendence basis x_1, \dots, x_n , in particular, dx_i form a basis. Then the elements $d(t_i + ax_i) = dt_i + adx_i$ form a basis for all but finitely many values of $a \in k$. In particular, we can choose $a \in k$ such that $d(t_i + ax_i)$ form a basis and $|ax_i| < 1$ for any i , and then $t'_i = t_i + ax_i$ are as required.

Note that the induced valuation on $l = k(t_1, \dots, t_n)$ is Gauss, and so l° is a localization of $k^\circ[t_1, \dots, t_n]$. Since $l^\circ \hookrightarrow K^\circ$ is a finite lci homomorphism, there exists a representation $K^\circ = l^\circ[s_1, \dots, s_m]/(f_1, \dots, f_m)$ and then K° is a localization of $k^\circ[\underline{t}, \underline{s}]/(f)$. By [28, 4.1.4], the canonical module ω_{K°/k° is generated by $\Delta^{-1}\phi$, where $\phi = dt_1 \wedge \dots \wedge dt_n$ and $\Delta = \det\left(\frac{\partial f_i}{\partial s_j}\right)$. To describe $\|\cdot\|_{\text{wt}}$ it suffices to compute the norm of ϕ . The definitions of [28, 4.2.3–4.2.5] introduce a log-norm that we call *weight*: $\text{wt}(\phi) = \frac{v_k(\Delta)+1}{e}$, where $e = e_{K/k}$ and $v_k : k^\times \rightarrow \mathbf{Z}$ is the additive valuation of k (the weight function wt_ϕ considered in loc.cit. is a function of a point $x \in X$ where K appears as the residue field of x ; it is the logarithmic analogue of the function $\|\phi\|$ on X). So, we define the *weight norm* by $\|\phi\|_{\text{wt}} = |\Delta\pi_K|$, where π_K is a uniformizer of K . (The factor $1/e$ is only needed in the additive setting to make the group of values of K equal to $e^{-1}\mathbf{Z}$, so that v_k agrees with v_{K° .) The weight norm on $\omega_{K/k}^{\otimes m}$ is defined via $\|\phi^{\otimes m}\|_{\text{wt}^{\otimes m}} = (\|\phi\|_{\text{wt}})^m$.

Note that Ω_{K°/l° is the cokernel of the map

$$\bigoplus_{i=1}^m f_i K^\circ \xrightarrow{d} \bigoplus_{j=1}^m K^\circ ds_j,$$

where $df_i = \sum_{j=1}^n \frac{\partial f_i}{\partial s_j} ds_j$. It follows that $\Delta = \text{cont}(\Omega_{K^\circ/l^\circ}) = \delta_{K/l}$.

8.3.2 Comparison

The norms $\| \cdot \|_\omega$ and $\| \cdot \|_{\text{wt}}$ on the one-dimensional vector space $\omega_{K/k}$ differ by a factor, so to compare them it suffices to evaluate the Kähler seminorm at ϕ . Since $|t_i| = 1$, Theorem 5.7.8 implies that dt_1, \dots, dt_n is a basis of the K° -module $\Omega_{I^\circ/k^\circ}^{\log} \otimes_{I^\circ} K^\circ$. The homomorphism $\psi_{K^\circ/I^\circ/k^\circ}^{\log}$ is an almost embedding by Lemma 5.4.5, and since its source is torsion free (even free), it is injective and we obtain an exact sequence

$$0 \rightarrow \Omega_{I^\circ/k^\circ}^{\log} \otimes_{I^\circ} K^\circ \rightarrow \Omega_{K^\circ/k^\circ}^{\log} \rightarrow \Omega_{K^\circ/I^\circ}^{\log} \rightarrow 0.$$

Note that

$$\left(\Omega_{K^\circ/k^\circ}^{\log} \right)_{\text{tf}} / \left(\Omega_{I^\circ/k^\circ}^{\log} \otimes_{I^\circ} K^\circ \right) = \Omega_{K^\circ/I^\circ}^{\log} / \left(\Omega_{K^\circ/k^\circ}^{\log} \right)_{\text{tor}}$$

and denote this module by M .

By Theorem 5.1.8, $(\Omega_{K^\circ/k^\circ}^{\log})_{\text{tf}}$ is an almost unit ball of $\| \cdot \|_{\omega, K/k}$. Since ϕ is a basis of $\det(\Omega_{I^\circ/k^\circ}^{\log} \otimes_{I^\circ} K^\circ)$ we have that

$$\| \phi \|_{\omega, K/k} = \left[\left(\Omega_{K^\circ/k^\circ}^{\log} \right)_{\text{tf}} : \left(\Omega_{I^\circ/k^\circ}^{\log} \otimes_{I^\circ} K^\circ \right) \right]^{-1},$$

and then Lemma 2.6.4 implies that $\| \phi \|_{\omega, K/k} = \text{cont}(M)$. By Theorem 2.6.7

$$\text{cont}(M) = \text{cont} \left(\Omega_{K^\circ/I^\circ}^{\log} \right) / \text{cont} \left(\left(\Omega_{K^\circ/k^\circ}^{\log} \right)_{\text{tor}} \right) = \delta_{K/I}^{\log} / \delta_{K/k}^{\log},$$

and we obtain that $\| \phi \|_\omega = \delta_{K/I}^{\log} / \delta_{K/k}^{\log}$. Since $\delta_{K/I}^{\log} = \delta_{K/I} | \pi_K \pi_I^{-1} |$ and $| \pi_I | = | \pi_k |$, the equality rewrites as $\| \phi \|_\omega = (\delta_{K/k}^{\log})^{-1} \Delta | \pi_K \pi_k^{-1} |$. Thus, $\| \cdot \|_{\text{wt}} = | \pi_k | \delta_{K/k}^{\log} \| \cdot \|_\omega$ and twisting by m we obtain the following comparison result.

Theorem 8.3.3. *If k is discretely valued, X is quasi-smooth and $x \in X$ is a divisorial point, i.e. a monomial point with discretely valued $K = \mathcal{H}(x)$, then the Kähler and the weight norms on m -canonical forms are related by*

$$\| \cdot \|_{\text{wt} \otimes m} = | \pi_k |^m \left(\delta_{K/k}^{\log} \right)^m \| \cdot \|_{\omega \otimes m}.$$

Remark 8.3.4. (i) If X is the analytification of a smooth k -variety, Mustařă and Nicaise extend the weight norms $\| \cdot \|_{\text{wt}, x}$ to a *weight seminorm* $\| \cdot \|_{X, \text{wt}}$ on the whole X by semicontinuity, i.e. $\| \cdot \|_{X, \text{wt}}$ is the minimal seminorm that extends the family $\{ \| \cdot \|_{\text{wt}, x} \}_{x \in X^{\text{div}}}$. If $\text{char}(\tilde{k}) = 0$, then $\| \cdot \|_{X, \omega \otimes m}$ is an X^{div} -seminorm by Corollary 8.2.10, hence Theorem 8.3.3 implies that $\| \cdot \|_{X, \text{wt} \otimes m} = | \pi_k |^m \| \cdot \|_{X, \omega \otimes m}$. If $\text{char}(\tilde{k}) > 0$ then it is easy to see that $X = \mathbf{A}_k^1$ contains divisorial points x with $\delta_{\mathcal{H}(x)/k}^{\log} < 1$. Hence the seminorms differ already for \mathbf{A}_k^1 .

- (ii) The constant factor $|\pi_k|$ in the formula for $\|\cdot\|_{\text{wt}}$ is analogous to the -1 shift in [28, 4.5.3], while the log different factor is rather subtle (whenever $\text{char}(\tilde{k}) > 0$). It seems very probable that for any K and quasi-smooth X , the function $\delta^{\log}(x) = \delta_{\mathcal{H}(x)/k}^{\log}$ is upper semicontinuous on X . In particular, the seminorms $\|\cdot\|_{\text{wt},x} = |\pi_k| \delta_{\mathcal{H}(x)/k}^{\log} \|\cdot\|_{\omega,x}$ should define an analytic seminorm $\|\cdot\|_{X,\text{wt}}$ on ω_X . Also, I expect that $\|\cdot\|_{X,\text{wt}}$ (as well as $\|\cdot\|_{X,\omega^{\otimes m}}$) is an X^{div} -seminorm, and hence it coincides with the weight seminorm of Mustața–Nicaise in the situation they considered.

9 The Topological Realization of X_G

Let X be an analytic space. In Sect. 9 we study the topological space $|X_G|$ associated with X_G and prove Theorem 9.4.8 that was used earlier in the paper. The section is independent of the rest of the paper, so there is no cycle reasoning here.

9.1 Topological Realization of X_G

9.1.1 Prime Filters

Let us recall the definition of completely prime filters on G -topological spaces (e.g., see [42, p. 83]). Let $p = \{U_i\}$ be a set of analytic domains of X then

- (1) p is *proper* if $\emptyset \notin p$ and $X \in p$.
- (2) p is *saturated* if for any analytic domains $U \subseteq V$ with $U \in p$ also $V \in p$.
- (3) p is *filtered* if for any $U, V \in p$ also $U \cap V \in p$.
- (4) p is *completely prime* if for any admissible covering $U = \cup_i U_i$ with $U \in p$ at least one U_i is in p .

We say that p is a *filter* if it is proper, saturated, and filtered. Note that these three conditions are purely set-theoretic, while the complete primality condition involves the G -topology.

Remark 9.1.2. (i) A filter is called *prime* if it satisfies (4) for finite admissible coverings. Van der Put and Schneider considered in [42] prime filters of *quasi-compact* analytic domains. In this case, primality and complete primality are equivalent. Moreover, any finite covering of a quasi-compact domain by quasi-compact domains is admissible, hence primality reduces to a set-theoretical condition.

- (ii) Our definition deals with arbitrary analytic domains. In this case, prime filters do not form an interesting class and one has to work with completely prime ones.

9.1.3 The Space $|X_G|$

By a *point* x of X_G we mean a completely prime filter $\{U_i\}$ of analytic domains of X . Intuitively, this is the prime filter of all analytic domains “containing” x . We denote by $|X_G|$ the set of all points of X_G . For any analytic domain $U \subseteq X$ saturation of a filter of U in X induces an embedding $|U_G| \hookrightarrow |X_G|$.

We provide $|X_G|$ with the topology whose base is formed by all sets of the form $|U_G|$. Obviously, any sheaf \mathcal{F} on X_G extends to a sheaf \mathcal{F}' on $|X_G|$ by setting $\mathcal{F}'(|U_G|) = \mathcal{F}(U)$ and sheafifying, so we obtain a functor $\alpha_X : X_G^\sim \rightarrow |X_G|^\sim$ between the associated topoi, where, as in [34], given a site \mathcal{C} we denote by \mathcal{C}^\sim the topos of sheaves of sets on \mathcal{C} . The stalk of \mathcal{F}' at $x \in |X_G|$ is simply $\text{colim}_{U \in x} \mathcal{F}(U)$. For shortness, we will denote this stalk as \mathcal{F}_x .

Remark 9.1.4. We refer to [36, Tag:00Y3] for the definition of points of a general site. It is easy to see that for G -topological spaces this definition agrees with our definition in terms of completely prime filters.

9.1.5 Abundance of Points

Since any point of X possesses a compact neighborhood and any compact analytic space is quasi-compact in the G -topology, the site of X is locally coherent in the sense of [35, VI.2.3]. Therefore, X_G^\sim has enough points by Deligne’s theorem, see [35, VI.9.0].

Theorem 9.1.6. *For any k -analytic space X the topological space $|X_G|$ is sober and the functor $\alpha_X : X_G^\sim \rightarrow |X_G|^\sim$ is an equivalence of categories.*

Proof. In view of [34, IV.7.1.9] and [34, VI.7.1.6], it suffices to show that X_G^\sim is generated by subsheaves of the final sheaf 1_X . For any analytic subdomain $U \subseteq X$, let $\mathcal{L}_U \subseteq 1_X$ denote the extension of 1_U , i.e. $\mathcal{L}_U(V) = \{1\}$ if $V \subseteq U$ and $\mathcal{L}_U(V) = \emptyset$ otherwise. If \mathcal{F} is a sheaf on X then any section $s \in \mathcal{F}(U)$ induces a morphism $\mathcal{L}_U \rightarrow \mathcal{F}$. In particular, we obtain an epimorphism $\phi : \coprod_{s,U} \mathcal{L}_U \rightarrow \mathcal{F}$, where U runs over all analytic subdomains and s runs over $\mathcal{F}(U)$. □

9.1.7 Ultrafilters

Our next aim is to classify the points of X_G and we start with those corresponding to the maximal completely prime filters.

Lemma 9.1.8. *The completely prime filter \mathcal{P}_x of all analytic domains containing a point $x \in X$ is maximal, and any maximal completely prime filter is of the form \mathcal{P}_x . In particular, we obtain an embedding of sets $X \hookrightarrow |X_G|$ whose image consists of all points that have no non-trivial generizations.*

Proof. If \mathcal{P}_x is not maximal then it can be increased to a larger completely prime filter \mathcal{P} . Fix an analytic domain $U \in \mathcal{P} \setminus \mathcal{P}_x$. Take an affinoid domain V containing x , then $W = U \cap V$ lies in \mathcal{P} . Choose an admissible covering of W by affinoid domains W_i . Then at least some $W' = W_i$ lies in \mathcal{P} . Since W' is a compact domain in V not containing x , there exists a neighborhood V' of x in V such that $V' \cap W' = \emptyset$. Since $V' \in \mathcal{P}_x \subset \mathcal{P}$, this contradicts \mathcal{P} being a filter, so \mathcal{P}_x is maximal.

Assume, now, that \mathcal{P} is a maximal completely prime filter. Since X possesses an admissible covering $X = \cup_i V_i$ by affinoid domains we can fix $V = V_i \in \mathcal{P}$. Assume that \mathcal{P} is not of the form \mathcal{P}_x with $x \in V$. By the maximality of \mathcal{P} , for any $x \in V$ we have that $\mathcal{P} \not\subseteq \mathcal{P}_x$ and hence there exists an affinoid domain $V_x \subset V$ with $x \notin V_x$. Then $\cap_{x \in V} V_x = \emptyset$, and hence already the intersection of finitely many sets V_x is empty. This contradicts \mathcal{P} being a filter. \square

In the sequel we will freely consider X as a subset of $|X_G|$.

Remark 9.1.9. One may wonder whether $X \hookrightarrow X_G$ is a topological embedding. Clearly, this may make sense only for the G -topology of X since any analytic domain $U \subseteq X$ is the preimage of the open subset U_G of X_G . Nevertheless, even for the G -topology the answer is negative simply because X is not a topological space for the G -topology. Moreover, there exist analytic domains U and V such that $U \cup V$ is not an analytic domain (e.g., the closed polydisc of radii $(1, 2)$ and the open polydisc of radii $(2, 1)$). So, $U_G \cup V_G$ is open in X_G but its restriction to X is not an analytic domain.

Corollary 9.1.10. *Any point $z \in |X_G|$ possesses a unique generization $\tau(z)$ lying in X .*

Proof. We should prove that any completely prime filter is contained in a single filter of the form \mathcal{P}_x . One such \mathcal{P}_x exists by Lemma 9.1.8. Assume that \mathcal{P} is contained in \mathcal{P}_x and \mathcal{P}_y with $x \neq y$. Choose any $V \in \mathcal{P}$, then $V_x = X \setminus \{y\}$ and $V_y = V \setminus \{x\}$ form an open and, hence, admissible covering of V . Thus, either V_x or V_y lies in \mathcal{P} and we obtain a contradiction. \square

9.1.11 The Retraction

By Corollary 9.1.10 we obtain a retraction $\tau_X : |X_G| \rightarrow X$ given by $z \mapsto \tau(z)$. For each $x \in X$, the fiber $\tau_X^{-1}(x)$ is the set of all specializations of x in $|X_G|$. Thus, $\tau_X^{-1}(x)$ is the closure of x in $|X_G|$ and we will also denote it $\bar{x}_{X,G}$.

Theorem 9.1.12. *Let X be a k -analytic space, $U \subseteq X$ a k -analytic subdomain, and $x \in U$ a point. Then,*

- (i) $\tau_X : |X_G| \rightarrow X$ is a topological quotient map and X is the maximal locally Hausdorff quotient of $|X_G|$.
- (ii) U is a neighborhood of x if and only if the inclusion $\bar{x}_{U,G} \subseteq \bar{x}_{X,G}$ is an equality. In particular, U is open if and only if the inclusion $|U_G| \subseteq \tau_X^{-1}(U)$ is an equality.

Proof. We start with (ii). For an analytic domain $V \subseteq X$ with $x \in V$ let (V, x) denote the germ of V at x . Define a presheaf of abelian groups on $|X_G|$ as follows: $F(V)$ is either 0 or \mathbf{Z} , and the second case takes place if and only if $x \in V$ and (V, x) is not contained in (U, x) , i.e. for any neighborhood W of x one has that $W \cap V \not\subseteq W \cap U$. The restriction maps are either identities or the map $\mathbf{Z} \rightarrow 0$. If V_1, \dots, V_n are domains containing x and satisfying $\cup_{i=1}^n (V_i, x) = (V, x)$, then $(V, x) \not\subseteq (U, x)$ if and only if $(V_i, x) \not\subseteq (U, x)$ for some i . It follows that the presheaf F is separated and hence the sheafification map $F \rightarrow \mathcal{F} = \alpha F$ is injective by [34, II.3.2]. In particular, $\mathcal{F} = 0$ if and only if $F = 0$. Obviously, $F = 0$ if and only if U is a neighborhood of x .

On the other hand, the stalk of \mathcal{F} at a point $z \in |X_G|$ is given by $\mathcal{F}_z = \text{colim}_{z \in W} F(W)$, in particular, \mathcal{F}_z is either 0 or \mathbf{Z} . The second possibility holds if and only if for any W with $z \in W$ one has that $x \in W$ and (W, x) is not contained in (U, x) . The first condition means that $\tau_X(z) = x$ and then the second condition holds if and only if $z \notin |U_G|$, i.e. $z \notin |U_G| \cap \bar{x}_{X,G} = \bar{x}_{U,G}$. Thus, \mathcal{F} has non-zero stalks if and only if the inclusion $\bar{x}_{U,G} \subseteq \bar{x}_{X,G}$ is not equality, and hence $\mathcal{F} = 0$ if and only if $\bar{x}_{U,G} = \bar{x}_{X,G}$. Combining this with the conclusion of the above paragraph we obtain (ii).

Let us prove (i). If $U \subseteq X$ is open then $\tau_X^{-1}(U) = |U_G|$ by (ii). Thus, $\tau_X^{-1}(U)$ is open in $|X_G|$ and we obtain that τ_X is continuous.

To prove that τ_X is a topological quotient map, assume that $U \subseteq X$ is not open and let us prove that $\tau_X^{-1}(U)$ is not open. Choose a point $x \in U$ not lying in the interior of U . By (ii) there exists a point $z \in \bar{x}_{X,G} \setminus \bar{x}_{U,G}$, in particular, $z \notin |U_G|$ and $z \in \tau_X^{-1}(U)$. We claim that $\tau_X^{-1}(U)$ is not a neighborhood of z . Assume to the contrary that z lies in the interior of $\tau_X^{-1}(U)$. Then there exists an analytic domain $W \subseteq X$ such that $z \in |W_G| \subseteq \tau_X^{-1}(U)$. Since $z \notin |U_G|$, we have that $W \not\subseteq U$. So there exists a point $y \in W \setminus U$, and observing that $y \in |W_G|$ and $y \notin \tau_X^{-1}(U)$ we obtain a contradiction. Thus, τ_X is a topological quotient map.

Finally, X is the maximal locally Hausdorff quotient of $|X_G|$ because any locally Hausdorff quotient $|X_G| \rightarrow Z$ should identify each point $x \in X$ with any of its specialization $z \in \bar{x}_{X,G}$. □

9.1.13 Notation X_G

Starting from this point we will not distinguish the site X_G and the topological space $|X_G|$. In particular, we will usually write $x \in X_G$ instead of $x \in |X_G|$.

9.1.14 Non-analytic Points

Our next aim is to describe the *non-analytic* or infinitesimal points of X_G , i.e. the points of $X_G \setminus X$. For this we have to recall some results about reductions of germs.

9.1.15 Germ Reduction

By $\widetilde{\mathcal{A}}_H$ we denote the H -graded reduction $\bigoplus_{h \in H} \widetilde{\mathcal{A}}_h$. In [10, Sect. 8] and [15, Sect. 1.5], to any germ (X, x) of an H -strict analytic space at a point one associates an H -graded reduction $(\widetilde{X}, \widetilde{x})_H$, which is a H -graded Riemann–Zariski space associated with the extension of H -graded fields $\widetilde{\mathcal{H}}(x)_H/\widetilde{k}_H$. In particular, any point $z \in (\widetilde{X}, \widetilde{x})_H$ induces a graded valuation on $\widetilde{\mathcal{H}}(x)_H$ and if (X, x) is separated then z is also determined by this valuation.

Theorem 9.1.16. *Let X be a H -strict k -analytic space and $x \in X$ a point. Then the closure of x in X_G is canonically homeomorphic to $(\widetilde{X}, \widetilde{x})_H$.*

Proof. Let L denote the set-theoretical lattice of subdomains $(U, x) \subseteq (X, x)$. Clearly, the closure of x is a sober topological space and L is its topology base. On the other hand, subdomains of $(\widetilde{X}, \widetilde{x})_H$ are in a one-to-one correspondence with quasi-compact open subspaces of $(\widetilde{X}, \widetilde{x})_H$ by [37, Theorem 4.5] and [10, Theorem 8.5]. So, L is also the lattice of a topology base of the sober topological space $(\widetilde{X}, \widetilde{x})_H$. Since, a sober topological space is determined by such a lattice (it can be reconstructed as points of the corresponding topos), we obtain the asserted homeomorphism. □

9.2 Stalks of \mathcal{O}_{X_G} and $\mathcal{O}_{X_G}^\circ$

Our next aim is to describe the stalks of the sheaves \mathcal{O}_{X_G} and $\mathcal{O}_{X_G}^\circ$ at non-analytic points. In particular, this will lead to an explicit description of the homeomorphism from Theorem 9.1.16.

9.2.1 Spectral Seminorm

We provide each stalk $\mathcal{O}_{X_G, x}$ with the stalk $|\cdot|_x$ of the spectral seminorm, i.e. $|s|_x = \inf_{x \in U_G} |s|_U$, where $|\cdot|_U$ is the spectral seminorm of U .

Lemma 9.2.2. *Let X be an analytic space and let $x \in X_G$ be a point. Then,*

- (i) *The seminorm $|\cdot|_x$ is a semivaluation and $\mathcal{O}_{X_G, x}$ is a local ring whose maximal ideal is the kernel of $|\cdot|_x$.*
- (ii) *If $y \in X_G$ generizes x , then the generization homomorphism $\phi_{x,y} : \mathcal{O}_{X_G, x} \rightarrow \mathcal{O}_{X_G, y}$ is an isometry with respect to $|\cdot|_x$ and $|\cdot|_y$. In particular, ϕ is local.*

For the sake of comparison we note that in the case of schemes any non-trivial generization is not local.

Proof. Let $z = \tau_X(x)$ be the maximal generization of x . We claim that $|s|_x = |s|_z$ for any $s \in \mathcal{O}_{X_G, x}$. If U is an analytic domain with $x \in U_G$ then $z \in U$ and so $|s|_x \geq |s|_z$.

Conversely, set $r = |s|_z$ and note that $X_\varepsilon = X\{|s| \leq r + \varepsilon\}$ is a neighborhood of z for any $\varepsilon > 0$. Since $x \in (X_\varepsilon)_G$ we obtain that $|s|_x \leq r + \varepsilon$ for any ε , and so $|s|_x \leq |s|_z$. In other words, $\phi_{x,z}$ is an isometry. In the same way, $\phi_{y,z}$ is an isometry, and therefore $\phi_{x,y}$ is an isometry.

It remains to prove (i). Since $\phi_{x,z}$ is an isometry and $|\cdot|_z$ is multiplicative, it follows that $|\cdot|_y$ is multiplicative. If $|s|_x = 0$ for $s \in \mathcal{O}_{X_G,x}$, then s is not invertible. Conversely, if $|s|_x = r > 0$, then $|s|_z > 0$ and hence there exists a neighborhood U of z such that $s \in \mathcal{O}_{X_G}(U)^\times$. Since $x \in U_G$, we obtain that s is invertible. \square

9.2.3 Residue Fields

Once we know that $\mathcal{O}_{X_G,x}$ is local, we denote its maximal ideal by $m_{G,x}$. The residue field will be denoted $\kappa_G(x) = \mathcal{O}_{X_G,x}/m_{G,x}$. Since $m_{G,x}$ is the kernel of $|\cdot|_x$, the residue field acquires a real valuation and we denote its completion by $\mathcal{H}(x)$. This extends the notation of Sect. 3.3.5 to non-analytic points.

9.2.4 Generization Homomorphisms

Recall that by Lemma 9.2.2(ii), any generization homomorphism $\phi_{x,y} : \mathcal{O}_{X_G,y} \rightarrow \mathcal{O}_{X_G,x}$ induces an embedding of real-valued fields $\kappa_G(y) \hookrightarrow \kappa_G(x)$.

Lemma 9.2.5. *If $y \in X_G$ generizes $x \in X_G$, then the induced embedding $\kappa_G(x) \hookrightarrow \kappa_G(y)$ has dense image and so $\mathcal{H}(x) = \mathcal{H}(y)$.*

Proof. Note that we can replace X with an analytic domain X' such that $x \in X'_G$. In particular, we can assume that X is affinoid. Let $z \in X$ be the maximal generization of y and x . The local embeddings $\mathcal{O}_{X,z} \hookrightarrow \mathcal{O}_{X_G,x} \hookrightarrow \mathcal{O}_{X_G,y} \hookrightarrow \mathcal{O}_{X_G,z}$ induce embeddings of the residue fields $\kappa(z) \hookrightarrow \kappa_G(x) \hookrightarrow \kappa_G(y) \hookrightarrow \kappa_G(z)$. It remains to use that $\kappa(z) \hookrightarrow \kappa_G(z)$ has a dense image. \square

9.2.6 Stalks of $\mathcal{O}_{X_G}^\circ$

As in [39, Sect. 2.1], by a semivaluation ring A with semifraction ring B we mean the following datum: a local ring (B, m) and a subring $A \subseteq B$ such that $m \subset A$ and A/m is a valuation ring of $k = B/m$. Such a datum defines an equivalence class of semivaluations $v : B \rightarrow \Gamma \cup \{0\}$ whose kernel is m , and $A = v^{-1}(\Gamma_{\leq 1} \cup \{0\})$ is the ring of integers of v .

Lemma 9.2.7. *For any $x \in X_G$, the stalk $\mathcal{O}_{X_G,x}^\circ$ is a semivaluation ring with semifraction ring $\mathcal{O}_{X_G,x}$. In particular, $\mathcal{O}_{X_G,x}^\circ/m_{G,x}$ is a valuation ring of $\kappa_G(x)$.*

Proof. If $s \in m_{G,x}$, then $|s|_x = 0$ and hence $|s| < 1$ in a neighborhood of x . In particular, $s \in \mathcal{O}_{X_G,x}^\circ$. It remains to show that if $u \in \mathcal{O}_{X_G,x}^\circ$ then either u or u^{-1}

lies in $\mathcal{O}_{X_G,x}^\circ$. Shrinking X we can assume that $u \in \mathcal{O}_{X_G}(X)^\times$. Then X is the union of $Y = X\{u\}$ and $Z = X\{u^{-1}\}$, hence x lies in either Y_G or Z_G , and then $u \in \mathcal{O}_{X_G,x}^\circ$ or $u^{-1} \in \mathcal{O}_{X_G,x}^\circ$, respectively. \square

9.2.8 Valuation v_x

Let v_x denote both the semivaluation induced by $\mathcal{O}_{X_G,x}^\circ$ on $\mathcal{O}_{X_G,x}$ and the valuations induced on $\kappa_G(x)$ and $\mathcal{H}(x)$. If $x \in X$, then an element $s \in \mathcal{O}_{X_G,x}$ lies in $\mathcal{O}_{X_G,x}^\circ$ if and only if $|s|_x \leq 1$. Thus, $\mathcal{O}_{X_G,x}^\circ/m_{G,x} = \mathcal{H}(x)^\circ$, i.e. v_x is the standard real valuation of $\mathcal{H}(x)$. If x is arbitrary, we only have an inclusion $\mathcal{O}_{X_G,x}^\circ/m_{G,x} \hookrightarrow \mathcal{H}(x)^\circ$, so v_x is composed from the real valuation of $\mathcal{H}(x)$ and the residue valuation \tilde{v}_x on $\widetilde{\mathcal{H}(x)}$.

The following remark clarifies the relation between these valuations and germ reductions. It will not be used in the sequel so we just formulate the results.

Remark 9.2.9. (i) The construction of \tilde{v}_x can be extended by associating with x a graded valuation ring of $\widetilde{\mathcal{H}(x)}_H$ (see Sect. 9.1.15) whose component in degree 1 is $\mathcal{O}_{X_G,x}^\circ/m_{G,x}$, and the argument is essentially the same. Set $\mathcal{O} = \mathcal{O}_{X_G}$ for shortness and let \mathcal{O}_r° and $\mathcal{O}_r^{\circ\circ}$ be the subsheaves of \mathcal{O} whose sections on U satisfy $|s|_U \leq r$ and $|s|_U < r$, respectively. Then $A_x = \bigoplus_{h \in H} (\mathcal{O}_h^\circ)_x / (\mathcal{O}_h^{\circ\circ})_x$ is a graded valuation ring of $\widetilde{\mathcal{H}(x)}_H = \bigoplus_{h \in H} (\mathcal{O}_x)_h^\circ / (\mathcal{O}_x)_h^{\circ\circ}$ that coincides with the graded valuation ring induced by the image of x under the homeomorphism $\tau_x^{-1}(y) = (\widetilde{X}, y)_H$ from Theorem 9.1.16, where $y \in X$ is the maximal generalization of x . In particular, $x \in X$ if and only if $A_x = \widetilde{\mathcal{H}(x)}_H$, i.e. the graded valuation is trivial.

(ii) Assume that $H = \sqrt{|k^\times|}$ and hence X_G is the usual strictly analytic G-topology. Then the H -graded reduction coincides (as a topological space) with the ungraded reduction (\widetilde{X}, x) because taking the degree-1 components provides a one-to-one correspondence between graded valuation \tilde{k}_H -rings in $\widetilde{\mathcal{H}(x)}_H$ and valuation \tilde{k} -rings in $\widetilde{\mathcal{H}(x)}$. In particular, x is determined by its maximal generalization y and a point of (\widetilde{X}, y) , or, that is equivalent, x is determined by the valuation v_x . This implies that X_G coincides with the Huber adic space X^{ad} corresponding to X , and $x \in X$ if and only if v_x is of height one. Furthermore, if $H = \mathbf{R}_{>0}$, then X_G coincides with the so-called reified adic space introduced by Kedlaya in [24], and (perhaps) the case of a general H will be set in detail in [16].

9.3 Topological Realization of PL Spaces and Skeletons

9.3.1 The n -Dimensional Affine R_S -PL Space

Consider the R_S -PL space $A = \mathbf{R}_{>0}^n$ with coordinates t_1, \dots, t_n . It is provided with the G -topology A_G of R_S -PL subspaces. Recall that $U \subseteq A$ is an R_S -PL subspace if the R_S -polytopes contained in U form a quasi-net (hence also a net) of U , and a covering $U = \cup_i U_i$ by R_S -PL subspaces is admissible if $\{U_i\}_i$ is a quasi-net of U .

For any polytope $P \subset A$ the topological realization $|P_G|$ is a quasi-compact topological space. As in the case of analytic spaces (see Sect. 9.1.5), Deligne’s theorem implies that A_G has enough points, and, moreover, $(A_G)^\sim$ is equivalent to $|A_G|^\sim$. For shortness, we will not distinguish A_G (resp. P_G) and its topological realization $|A_G|$ (resp. $|P_G|$).

9.3.2 G -Skeletons

If X is an analytic space with a \mathbf{Z}_H -PL subspace P , then the embedding $i : P \hookrightarrow X$ is continuous with respect to the G -topologies, see [7, Theorem 6.3.1]. Since the functor that associates with a sober topological space the lattice of its open subsets is fully faithful, this implies that i extends to a continuous embedding $i_G : P_G \hookrightarrow X_G$ and we say that P_G is a \mathbf{Z}_H -PL subspace of X_G .

Remark 9.3.3. (i) The main advantage of working with P_G is that it is a honest topological space, so one can use local arguments. One has to describe the new points but, as we will see, this is simple: points of P_G correspond to valuations on abelian groups, and points of $i_G(P_G)$ correspond to \underline{t} -monomial valuations.

(ii) It seems that the use of model theory in [16] is mainly needed for the same aim. One interprets P_G in terms of definable sets and types in the theory of ordered groups, and Deligne’s theorem on points of locally coherent sites is replaced with Gödel’s completeness theorem. This is not so surprising, since it is known that Gödel’s theorem and Deligne’s theorem are equivalent, when appropriately translated (for example, see [17]).

9.3.4 Monomiality

Notions of \underline{t} -monomial and generalized Gauss valuations naturally extend to general valuations. Namely, let L/l be an extension of valued fields and let (t_1, \dots, t_n) be a tuple of elements of L . We say that the valuation on $l(\underline{t})$ is a generalized Gauss valuation (with respect to l) if for any polynomial $a = \sum_{i \in \mathbf{N}^n} a_i \underline{t}^i \in l[\underline{t}]$ the equality $|a| = \max_i |a_i \underline{t}^i|$ holds. Such a valuation on $l(\underline{t})$ is uniquely determined by its restrictions onto l and the monoid $\underline{t}^{\mathbf{Z}} := \prod_{j=1}^n \underline{t}_j^{\mathbf{Z}}$. If, in addition, L is finite over the closure of $l(\underline{t})$ in L , then we say that L and its valuation are \underline{t} -monomial. The following result is proved in [16, Proposition 1.8.3].

Lemma 9.3.5. *Assume that X is an analytic space, $f : U \rightarrow \mathbf{G}_m^n$ is a monomial chart given by $t_1, \dots, t_n \in \mathcal{O}_{X_G}^\times(U)$, and $x \in X_G$ is a point whose maximal generization y is contained in the \mathbf{Z}_H -PL subspace $P = S(f)$. Let l denote the field $\widetilde{\mathcal{H}}(y)$ provided with the valuation \tilde{v}_x [see Remark 9.2.9(i)], and let $d = \text{tr.deg.}(l/\tilde{k})$. Assume that $|t_i|_x = 1$ for $1 \leq i \leq d$, and set $s_i = \tilde{t}_i$ for $1 \leq i \leq d$. Then $x \in P_G$ if and only if the extension l/\tilde{k} is s -monomial.*

Consider a monomial chart $f : U \rightarrow \mathbf{G}_m^n$ given by $u_1, \dots, u_n \in \mathcal{O}_{X_G}^\times(U)$. Let $g : U \rightarrow \mathbf{G}_m^n$ be given by $t_i = c_i \prod_{j=1}^n u_j^{l_{ij}}$, where $c_i \in k^\times$ and (l_{ij}) is an n -by- n integer matrix with non-zero determinant. Then it is easy to see that g is another monomial chart and $S(f) = S(g)$. Note also that if $x \in U_G$ is a point whose maximal generization y is monomial then choosing monomials t_i appropriately we can achieve that $|t_i|_x = 1$ for $1 \leq i \leq F_y = \text{tr.deg.}(\widetilde{\mathcal{H}}(x)/\tilde{k})$ (in terms of [16, Sect. 1.8], t_i are well presented at x). Therefore, Lemma 9.3.5 implies the following corollary.

Corollary 9.3.6. *Assume that P is a \mathbf{Z}_H -PL subspace of X and $x \in P_G$ is a point, and let l denote the residue field $\widetilde{\mathcal{H}}(x)$ with the valuation \tilde{v}_x . Then the extension l/\tilde{k} is Abhyankar.*

9.3.7 Structure Sheaf of A

In the remaining part of Sect. 9.3 we provide the promised valuation-theoretic description of the points of P_G . This will not be used in the sequel, so the reader can skip to Sect. 9.4.

Until the end of Sect. 9.3, we fix arbitrary R and S , and $L = S \oplus (\bigoplus_{i=1}^n t_i^R)$ denotes the group of R_S -monomial functions on A . We provide A with the sheaf \mathcal{O}_A° such that if P is an R_S -polyhedron then $\mathcal{O}_A^\circ(P)$ is the monoid of R_S -PL functions with values in $(0, 1]$.

9.3.8 Combinatorial Valuations

By an R_S -valuation on L we mean any homomorphism $|\cdot| : L \rightarrow \Gamma$ to an ordered group such that if $r \in R_+ = R \cap \mathbf{R}_{\geq 0}$ and $x \in L$ with $|x| \leq 1$ then $|x^r| \leq 1$, and the restriction of $|\cdot|$ onto $S \subseteq L$ is equivalent to the embedding $S \hookrightarrow \mathbf{R}_{>0}$. We say that a valuation is *bounded* if for any $x \in L$ there exists $s \in S$ with $|x| \leq |s|$.

9.3.9 Valuation Monoids

Analogously to ring valuations, the valuation is determined up to an equivalence by the *valuation monoid* L° consisting of all elements $x \in L$ with $|x| \leq 1$. It is bounded

if and only if $SL^\circ = L$. In addition, $(L^\circ)^{\text{gp}} = L$ and L° is an $(R_S)^\circ$ -monoid, i.e. it contains $S^\circ = S \cap (0, 1]$ and is closed under the action of R_+ . In this case, we say that L° is a *bounded valuation R_+ -monoid of L* .

9.3.10 Valuative Interpretation of Points

Similarly to the points of X_G , points of P_G admit a simple valuative-theoretic interpretation.

Theorem 9.3.11. *Let A and L be as above. Then for any point $x \in A_G$ the stalk $\mathcal{O}_{A_G, x}^\circ$ is a bounded valuation R_+ -monoid of L , and this establishes a one-to-one correspondence between points of P_G and bounded R_S -valuations on L .*

Proof. Choose an R_S -polytope P with $x \in P_G$. Clearly, $M = \mathcal{O}_{A_G, x}^\circ$ is an R_+ -monoid. Also, $M^{\text{gp}} = L = SM$ because these equalities hold already for the monoid $M' = \mathcal{O}_{A_G}^\circ(P)$. To prove that M is a valuation monoid it suffices to show that if $a \in L$ then either $a \in M$ or $a^{-1} \in M$, but this follows from the fact that $P = P\{a \leq 1\} \cup P\{a^{-1} \leq 1\}$.

It remains to show that any bounded valuation R_+ -monoid L° of L equals to $\mathcal{O}_{A_G, x}^\circ$ for a unique point x . Consider the set \mathcal{F} of all R_S -polytopes given by finitely many inequalities $a_i \leq 1$ with $a_i \in L^\circ$. Then \mathcal{F} is a completely prime filter and the stalk at the corresponding point x is L° . Uniqueness of \mathcal{F} is also clear from the construction. □

9.3.12 G-Skeleton of the Torus

Now let us assume that $R_S = \mathbf{Z}_H$ and so $L = H \oplus \underline{t}^{\mathbf{Z}}$. Set $\mathbf{T} = \mathbf{G}_m^n$, and consider the embedding $i : A \hookrightarrow \mathbf{T}$ sending \underline{s} to the generalized Gauss valuation $|\cdot|_{\underline{s}}$ and the retraction $r : \mathbf{T} \rightarrow A$ from Remark 7.2.2(ii). Both are continuous in the G-topology, see [7, Theorems 6.3.1 and 6.4.1], hence extend to continuous maps $i_G : A_G \hookrightarrow \mathbf{T}_G$ and $r_G : \mathbf{T}_G \rightarrow A_G$. Furthermore, the description of the maps i and r can be naturally extended to i_G and r_G . For simplicity, we explain this only in the case when $H = |k^\times|^\mathbf{Q}$, and the reduction is ungraded.

Given a point $x \in \mathbf{T}_G$ consider its maximal generalization $y \in \mathbf{T}$ and provide $\mathcal{H}(x)$ with the valuation v_x , see Remark 9.2.9. Then the restriction of v_x onto $|k^\times| \oplus \underline{t}^{\mathbf{Z}}$ is a bounded \mathbf{Z}_H -valuation, so we obtain a map r_G . Conversely, given a bounded \mathbf{Z}_H -valuation $\mu : L \rightarrow \Gamma$, we extend it to a generalized Gauss valuation on $k[\underline{t}]$ by the max formula $|\sum_i a_i \underline{t}^i|_\mu = \max_i |a_i|_\mu(\underline{t}^i)$. Since $|\cdot|_\mu$ is composed from a real valuation and a valuation on its residue field, we obtain a point of \mathbf{T}_G . Clearly, $r_G \circ i_G = \text{Id}$, so r_G is a retraction onto $S(\mathbf{T})_G := i_G(A)$.

9.4 Residually Unramified Monomial Charts

Our last goal is to prove a theorem on existence of residually unramified monomial charts that was used earlier in the paper.

9.4.1 Residual Unramifiedness at Non-analytic Points

Assume that $f : Y \rightarrow X$ is a morphism of k -analytic spaces, $y \in Y_G$, $x = f(y)$ and provide $\mathcal{H}(y)$ and $\mathcal{H}(x)$ with the valuations v_y and v_x , respectively (see Sect. 9.2.8). We say that f is *residually unramified at y* (resp. *residually tame at y*) if the extension of valued fields $(\mathcal{H}(y), v_y)/(\mathcal{H}(x), v_x)$ is finite and unramified (resp. tame). This extends the analogous notion from the case of analytic points to the whole Y_G .

Lemma 9.4.2. *Keep the above notation. Then f is residually unramified at y if and only if the extension of the real-valued fields $(\mathcal{H}(y), | \cdot |_y)/(\mathcal{H}(x), | \cdot |_x)$ and the extension of the valued fields $(\widetilde{\mathcal{H}}(y), \widetilde{v}_y)/(\widetilde{\mathcal{H}}(x), \widetilde{v}_x)$ are unramified.*

Proof. This follows from a criterion for an extension of composed valued fields to be unramified, see [38, Proposition 2.2.2]. \square

9.4.3 Generators of Unramified Extensions

Assume that L/K is a finite unramified extension of valued fields. We say that $u \in L$ is an *integral generator* of L over K if L° is a localization of $K^\circ[u]$. Note that in this case $g'(u)$ is invertible in L° , where $g(T)$ is the minimal polynomial of u over K .

Lemma 9.4.4. *Assume that L/K is a finite extension of valued fields. Then,*

- (i) *If L/K is unramified, then it possesses an integral generator.*
- (ii) *An element $u \in L$ is an integral generator if and only if $L = K[u]$, $u \in L^\circ$ and $g'(u) \in (L^\circ)^\times$, where g is the minimal polynomial of u over K .*

Proof. The first claim is a (simple) special case of Chevalley's theorem [19, IV₄, 18.4.6]. Only the inverse implication needs a proof in (ii), so let us establish it. Consider the subring $A = K^\circ[u, g'(u)]$ of L° . By [19, IV₄, 18.4.2(ii)], A is étale over K° . Therefore, A is a semilocal Prüfer ring (i.e., all its localizations are valuation rings) and it follows that any intermediate ring $A \subseteq R \subseteq \text{Frac}(A)$ is a localization of A . In particular, L° is a localization of A . \square

9.4.5 Residually Unramified Locus

Given a morphism $f : Y \rightarrow X$, by the *residually unramified locus* of f we mean the set of all points $y \in |Y_G|$ such that f is residually unramified at y .

Theorem 9.4.6. *Assume that $f : Y \rightarrow X$ is a quasi-étale morphism. Then the residually unramified locus of f is open in $|Y_G|$.*

Proof. Assume that f is residually unramified at $y \in Y_G$. We should prove that f is residually unramified in a neighborhood of y . Set $x = f(y)$, $L = (\mathcal{H}(y), v_y)$ and $K = (\mathcal{H}(x), v_x)$. Also, let $y_0 \in Y$ and $x_0 \in X$ be the maximal generalizations of y and x , respectively, and consider the real-valued fields $L_0 = \mathcal{H}(y_0)$ and $K_0 = \mathcal{H}(x_0)$.

Step 1. *We can assume that $X = \mathcal{M}(\mathcal{A})$ is affinoid and f is finite étale.*

The question is G -local at x and y hence we can assume that X and Y are affinoid. It follows easily from [5, Theorem 3.4.1] that replacing X and Y with neighborhoods of y_0 and x_0 , we can achieve that f factors as $Y \hookrightarrow \bar{Y} \rightarrow X$, where Y is an affinoid domain in \bar{Y} and $\bar{Y} \rightarrow X$ is finite étale (see [16, 3.12]). So, replacing Y with \bar{Y} we can assume that f is finite étale.

Note that X and Y are good. We will use the usual local rings \mathcal{O}_{X,x_0} , \mathcal{O}_{Y,y_0} and the residue fields $\kappa(x_0)$, $\kappa(y_0)$ in the sequel. Also, we will freely replace X with a neighborhood of x_0 when needed.

Step 2. *We can assume that $Y = \mathcal{M}(\mathcal{B})$, where $\mathcal{B} = \mathcal{A}[u]$ and $u(y)$ is an integral generator of L° over K° .* By our assumption, L°/K° is étale, hence by Lemma 9.4.4(i) there exists an integral generator v of L over K . We claim that using Lemma 9.4.4(ii) one can slightly move v achieving that $v \in \kappa(y_0)$. Indeed, L_0/K_0 is separable hence $K_0[v]$ is preserved under small deformations of u by Krasner’s lemma applied to L_0/K_0 . Similarly, if g_v is the minimal polynomial of v , then $g'_v(v)$ changes slightly under small deformations of v . In particular, a slight change of v preserves the reduction $\widetilde{g'_v(v)}$ and hence also the equality $\tilde{v}_y(g'_v(v)) = 1$, which is equivalent to the inclusion $g'_v(v) \in (L^\circ)^\times$.

In the sequel, $v \in \kappa(y_0)$. Since $\kappa(x_0)$ is henselian by [5, Theorem 2.3.3], it is separably closed in $\mathcal{H}(x)$ and therefore the minimal polynomial $g_v(T)$ of v lies in $\kappa(x_0)[T]$. Choose a monic lifting $G(T) \in \mathcal{O}_{X,x_0}[T]$ and shrink X around x_0 so that $G(T) \in \mathcal{A}[T]$.

Set $\mathcal{B} = \mathcal{A}[T]/(G(T))$ and $Y' = \mathcal{M}(\mathcal{B})$. Then $Y' \rightarrow X$ is a finite map of degree $d = \deg(G)$ and the preimage of x_0 is a single point y'_0 such that $\mathcal{H}(y'_0)/\mathcal{H}(x_0)$ is the extension L_0/K_0 . Therefore, $Y' \rightarrow X$ is étale at y'_0 and the maps of germs $(Y', y'_0) \rightarrow (X, x_0)$ and $(Y, y_0) \rightarrow (X, x_0)$ are isomorphic by [5, Theorem 3.4.1]. In particular, after shrinking X around x_0 the morphism $Y' \rightarrow X$ becomes étale, and then we can replace Y with Y' . Then the image $u \in \mathcal{B}$ of T is as required since $u(y) = v$.

Step 3. *The domain $U = Y\{G'(u)^{-1}\}$ is as required.* Clearly $y \in U$, so we should only check that for any $z \in U$ with $F = \mathcal{H}(z)$ and $E = \mathcal{H}(f(z))$, the extension F/E is unramified. Set $w = u(z)$, then $F = E[w]$ and the minimal polynomial $h(T)$ of w over E is a factor of $\overline{G}(T)$, where $\overline{G} = G(z)$ is obtained from G by evaluating its coefficients at z . Since $h(w) = 0$ and $\overline{G}'(w)$ is invertible, we obtain that $h'(w)$ is invertible. By Lemma 9.4.4(ii), w is an integral generator of F over E , in particular, F/E is unramified. □

9.4.7 Construction of Charts

Now we are in a position to prove the main result of Sect. 9.4.

Theorem 9.4.8. *Assume that k is algebraically closed. Let X be a k -analytic space and P a compact \mathbf{Z}_H -PL subspace of X . Then there exist finitely many residually unramified monomial charts $f_i : U_i \rightarrow \mathbf{G}_m^{n_i}$ such that $P = \cup_i S(f_i)$.*

Proof. First, we observe that it suffices to show that for any point $x \in P_G$ there exists a residually unramified monomial chart $f : U \rightarrow \mathbf{G}_m^n$ such that $x \in S(f)_G$. Indeed, intersecting this chart with a chart g such that $x \in S(g)_G$ and $S(g) \subseteq P$ we can also achieve that $S(f) \subseteq P$, and then the assertion follows from the quasi-compactness of P_G .

By Lemma 9.4.9 below, there exists a monomial chart such that $x \in S(f)_G$ and f is residually unramified at x . Applying Theorem 9.4.6 we can find an analytic domain $V \subseteq U$ such that $x \in V_G$ and the map $f|_V$ is residually unramified. Hence $f|_V$ is a required monomial chart and we are done. \square

Lemma 9.4.9. *Assume that X is an analytic space and $x \in X_G$ is a monomial point. Then there exists a monomial chart $f : U \rightarrow \mathbf{G}_m^n$ such that $x \in S(f)_G$ and f is residually monomial at x .*

Proof. The argument is similar to that in Corollary 7.2.6, but this time we should choose the transcendence basis of the residue field more carefully. Let $x_0 \in X$ be the maximal generalization of x and consider the real-valued field $L_0 = (\mathcal{H}(x_0), | \cdot |_x)$ and the valued fields $L = (\mathcal{H}(x), v_x)$ and $l = (\tilde{L}_0, \tilde{v}_x)$. Then x_0 is monomial and l/\tilde{k} is Abhyankar by Corollary 9.3.6. Set $E = E_{x_0}$ and $F = F_{x_0}$, and choose elements a_{F+1}, \dots, a_{F+E} such that their images in $|L_0^\times|/|k^\times|$ form a basis. Next, choose a transcendence basis b_1, \dots, b_F of l/\tilde{k} such that $l/\tilde{k}(b_1, \dots, b_F)$ is unramified; this is possible by [26, Theorem 1.3].

Remark 9.4.10. Existence of such a basis is the valuation-theoretic ingredient of local uniformization of Abhyankar valuations. It is an immediate consequence of the difficult theorem that l is stable, see [26, Theorem 1.1] or [40, Remark 2.1.3]. In fact, one can take b_1, \dots, b_F to be any basis such that $b_1, \dots, b_{\tilde{E}}$ is mapped to a basis of $|l^\times| \otimes \mathbf{Q}$, where $\tilde{E} = E_{l/\tilde{k}}$, and $b_{\tilde{E}+1}, \dots, b_F$ is mapped to a separable transcendence basis of \tilde{l}/\tilde{k} .

Choose any lifts $a_1, \dots, a_F \in \kappa_G(x)$ of b_1, \dots, b_F . We obtain elements a_1, \dots, a_n , where $n = E + F = \dim_{x_0}(X)$, and let $t_1, \dots, t_n \in \mathcal{O}_{X_G, x}$ be any lifts of a_1, \dots, a_n . Let U be such that t_i are defined on U and $x \in U_G$. Then \underline{t} induces a morphism $f : U \rightarrow \mathbf{G}_m^n$ and the fiber $f^{-1}(f(x_0))$ is zero-dimensional at x_0 by [15, Corollary 8.4.3]. It then follows from [12, Theorem 4.9] that replacing U by a Zariski open neighborhood of x_0 we can achieve that f has zero-dimensional fibers and hence is a monomial chart. By our construction, the induced valuations on $K = k(a_1, \dots, a_n)$ and $\tilde{K} = \tilde{k}(b_1, \dots, b_F)$ are generalized Gauss valuations. The first implies that $x_0 \in S(f)$, hence the second implies that $x \in S(f)_G$ by Lemma 9.3.5.

The residue field extension $l/\tilde{k}(b_1, \dots, b_F)$ is unramified and hence separable. In addition, L_0 is unramified over $\mathcal{H}(f(x_0)) = \overline{k(a_1, \dots, a_F)}$ because the fields are stable and have the same group of values and the extension of the residue fields is separable. Hence f is residually unramified at x by Lemma 9.4.2, and we are done. \square

Acknowledgements This work originated at Simons Symposium on Non-Archimedean and Tropical Geometry in 2013, where I learned about the work of Mustařa and Nicaise and found an analytic approach to metrize pluriforms. I want to thank Simons Foundation and the organizers of the Symposium, Matt Baker and Sam Payne. I owe much to the anonymous referees for pointing out numerous inaccuracies and making various suggestions on improving the quality and readability of the paper. Also, I am grateful to J. Nicaise for answering questions related to [28], to A. Ducros for showing his work with Thuillier [16] and answering various questions, and to V. Berkovich, I. Tyomkin, and Y. Varshavsky for useful discussions.

This work was supported by the European Union Seventh Framework Programme (FP7/2007–2013) under grant agreement 268182.

References

1. Abramovich, D., Karu, K.: Weak semistable reduction in characteristic 0. *Invent. Math.* **139**(2), 241–273 (2000). MR 1738451 (2001f:14021)
2. Ax, J.: Zeros of polynomials over local fields—the galois action. *J. Algebra* **15**, 417–428 (1970). MR 0263786 (41 #8386)
3. Ben-Bassat, O., Kremnizer, K.: Non-archimedean analytic geometry as relative algebraic geometry. ArXiv e-prints (2013). <http://arxiv.org/abs/1312.0338>
4. Berkovich, V.G.: Spectral theory and analytic geometry over non-Archimedean fields. *Mathematical Surveys and Monographs*, vol. 33. American Mathematical Society, Providence, RI (1990). MR 1070709 (91k:32038)
5. Berkovich, V.G.: Étale cohomology for non-Archimedean analytic spaces. *Inst. Hautes Études Sci. Publ. Math.* **78**, 5–161 (1993). MR 1259429 (95c:14017)
6. Berkovich, V.G.: Smooth p -adic analytic spaces are locally contractible. *Invent. Math.* **137**(1), 1–84 (1999). MR 1702143 (2000i:14028)
7. Berkovich, V.G.: Smooth p -adic analytic spaces are locally contractible. II. In: *Geometric Aspects of Dwork Theory*, vol. I, II, pp. 293–370. Walter de Gruyter, Berlin (2004). MR 2023293 (2005h:14057)
8. Bosch, S., Güntzer, U., Remmert, R.: *Non-Archimedean Analysis: A Systematic Approach to Rigid Analytic Geometry*. *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*, vol. 261. Springer, Berlin (1984). MR 746961 (86b:32031)
9. Chambert-Loir, A., Ducros, A.: Formes différentielles réelles et courants sur les espaces de Berkovich. ArXiv e-prints (2012). <http://arxiv.org/abs/1204.6277>
10. Conrad, B., Temkin, M.: Descent for non-archimedean analytic spaces <http://www.math.huji.ac.il/~temkin/papers/Descent.pdf>, (2007, in preparation)
11. Ducros, A.: Image réciproque du squelette par un morphisme entre espaces de Berkovich de même dimension. *Bull. Soc. Math. France* **131**(4), 483–506 (2003). MR 2044492 (2004m:14042)
12. Ducros, A.: Variation de la dimension relative en géométrie analytique p -adique. *Compos. Math.* **143**(6), 1511–1532 (2007). MR 2371379 (2008j:14046)

13. Ducros, A.: Espaces de Berkovich, polytopes, squelettes et théorie des modèles. *Confluentes Math.* **4**(4), 1250007, 57 (2012). MR 3020334
14. Ducros, A.: Toute forme modérément ramifiée d'un polydisque ouvert est triviale. *Math. Z.* **273**(1–2), 331–353 (2013). MR 3010163
15. Ducros, A.: Families of Berkovich spaces (2014, in preparation). <https://www.imj-prg.fr/~antoine.ducros/Families-Berko.pdf>
16. Ducros, A., Thuillier, A.: Squelettes et valuations monomiales (2014, in preparation)
17. Frot, B.: Gödel's completeness theorem and Deligne's theorem. *ArXiv e-prints* (2013). <http://arxiv.org/abs/1309.0389>
18. Gabber, O., Ramero, L.: Almost Ring Theory. *Lecture Notes in Mathematics*, vol. 1800. Springer, Berlin (2003). MR 2004652 (2004k:13027)
19. Grothendieck, A.: Éléments de géométrie algébrique. *Inst. Hautes Études Sci. Publ. Math.* **4,8,11,17,20,24,28,32** (1960–1967)
20. Hrushovski, E., Loeser, H.: Non-archimedean tame topology and stably dominated types. *ArXiv e-prints* (2012). <http://arxiv.org/abs/1009.0252>
21. Huber, R.: A generalization of formal schemes and rigid analytic varieties. *Math. Z.* **217**(4), 513–551 (1994). MR 1306024 (95k:14001)
22. Illusie, L.: *Complexe Cotangent et Déformations. I.* *Lecture Notes in Mathematics*, vol. 239. Springer, Berlin (1971)
23. Kato, K.: Logarithmic structures of Fontaine-Illusie. In: *Algebraic Analysis, Geometry, and Number Theory* (Baltimore, MD, 1988), pp. 191–224. Johns Hopkins University Press, Baltimore, MD (1989). MR 1463703 (99b:14020)
24. Kedlaya, K.: Reified valuations and adic spectra. *ArXiv e-prints* (2013). <http://arxiv.org/abs/1309.0574>
25. Kontsevich, M., Soibelman, Y.: Affine structures and non-Archimedean analytic spaces. In: *The Unity of Mathematics. Progress in Mathematics*, vol. 244, pp. 321–385. Birkhäuser Boston, MA (2006). MR 2181810 (2006j:14054)
26. Kuhlmann, F.-V.: Elimination of ramification I: the generalized stability theorem. *Trans. Amer. Math. Soc.* **362**(11), 5697–5727 (2010). MR 2661493 (2011e:12013)
27. Matsumura, H.: *Commutative Algebra*, 2nd edn. *Mathematics Lecture Note Series*, vol. 56. Benjamin/Cummings, Reading, MA (1980). MR 575344 (82i:13003)
28. Mustață, M., Nicaise, J.: Weight functions on non-archimedean analytic spaces and the Kontsevich-Soibelman skeleton. *ArXiv e-prints* (2013). <http://arxiv.org/abs/1212.6328>
29. Nicaise, J., Xu, C.: The essential skeleton of a degeneration of algebraic varieties. *ArXiv e-prints* (2013). <http://arxiv.org/abs/1307.4041>
30. Ogus, A.: Lectures on logarithmic algebraic geometry. math.berkeley.edu/~ogus/preprints/log_book/logbook.pdf
31. Olsson, M.C.: The logarithmic cotangent complex. *Math. Ann.* **333**(4), 859–931 (2005). MR 2195148 (2006j:14017)
32. Poineau, J.: Les espaces de Berkovich sont angéliques. *Bull. Soc. Math. France* **141**(2), 267–297 (2013). MR 3081557
33. Scholze, P.: Perfectoid spaces. *Publ. Math. Inst. Hautes Études Sci.* **116**, 245–313 (2012). MR 3090258
34. Théorie des topos et cohomologie étale des schémas. Tome 1: Théorie des topos, Springer-Verlag, Berlin, 1972, Séminaire de Géométrie Algébrique du Bois-Marie 1963–1964 (SGA 4), Dirigé par Michael Artin, Alexander Grothendieck, et Jean-Louis Verdier. Avec la collaboration de Nicolas Bourbaki, Pierre Deligne et Bernard Saint-Donat, *Lecture Notes in Mathematics*, vol. 269
35. Théorie des topos et cohomologie étale des schémas. Tome 2. Springer-Verlag, Berlin, 1972, Séminaire de Géométrie Algébrique du Bois-Marie 1963–1964 (SGA 4), Dirigé par Michael Artin, Alexander Grothendieck, et Jean-Louis Verdier. Avec la collaboration de Nicolas Bourbaki, Pierre Deligne et Bernard Saint-Donat, *Lecture Notes in Mathematics*, vol. 270
36. The Stacks Project Authors. *Stacks Project*. <http://stacks.math.columbia.edu>

37. Temkin, M.: On local properties of non-Archimedean analytic spaces. II. *Israel J. Math.* **140**, 1–27 (2004). MR 2054837 (2005c:14030)
38. Temkin, M.: Stable modification of relative curves. *J. Algebraic Geom.* **19**(4), 603–677 (2010). MR 2669727 (2011j:14064)
39. Temkin, M.: Relative Riemann-Zariski spaces. *Israel J. Math.* **185**, 1–42 (2011). MR 2837126
40. Temkin, M.: Inseparable local uniformization. *J. Algebra* **373**, 65–119 (2013). MR 2995017
41. Temkin, M.: Altered local uniformization of Berkovich spaces. ArXiv e-prints (2015). <http://arxiv.org/abs/1503.05729>
42. van der Put, M., Schneider, P.: Points and topologies in rigid geometry. *Math. Ann.* **302**(1), 81–103 (1995). MR 1329448 (96k:32070)

Skeletons and Fans of Logarithmic Structures

Dan Abramovich, Qile Chen, Steffen Marcus, Martin Ulirsch,
and Jonathan Wise

Abstract We survey a collection of closely related methods for generalizing fans of toric varieties, include skeletons, Kato fans, Artin fans, and polyhedral cone complexes, all of which apply in the wider context of logarithmic geometry. Under appropriate assumptions these structures are equivalent, but their different realizations have provided for surprisingly disparate uses. We highlight several current applications and suggest some future possibilities.

Keywords Logarithmic structures • Non-Archimedean geometry • Toric varieties • Tropical geometry • Polyhedral complexes • Algebraic stacks • Moduli spaces

D. Abramovich (✉)

Department of Mathematics, Brown University, Box 1917, Providence, RI 02912, USA
e-mail: abrmovic@math.brown.edu

Q. Chen

Department of Mathematics, Maloney Hall, Fifth Floor, Boston College,
Chestnut Hill, Boston, MA 02467-3806, USA
e-mail: Qile.Chen@bc.edu

S. Marcus

Department of Mathematics and Statistics, The College of New Jersey, Ewing, NJ 08628, USA
e-mail: marcuss@tcnj.edu

M. Ulirsch

Hausdorff Center for Mathematics, University of Bonn, 53115 Bonn, Germany
e-mail: ulirsch@math.uni-bonn.de

J. Wise

University of Colorado, Boulder, CO 80309-0395, USA
e-mail: jonathan.wise@math.colorado.edu

1 Introduction

1.1 Toric Varieties, Toroidal Embeddings, and Logarithmic Structures

Toric varieties were introduced in [21] and studied in many sources, see, for instance, [19, 20, 23, 35, 43, 44]. They are the quintessence of combinatorial algebraic geometry: there is a category of combinatorial objects called fans, which is equivalent to the category of toric varieties with torus-equivariant morphisms between them. Further, classical tropical geometry probes the combinatorial structure of subvarieties of toric varieties. We review this theory in utmost brevity in Sect. 2 below.

In [35] some of this picture was generalized to *toroidal embeddings*, especially *toroidal embeddings without self-intersections*, and their corresponding *polyhedral cone complexes* (Sect. 2.7). Following Kato [34] we argue that the correct generality is that of *fine and saturated logarithmic structures*. Working over perfect fields, toroidal embeddings can be identified as logarithmically regular varieties. Since logarithmic structures might not be familiar to the reader, we provide a brief review of the necessary definitions in Sect. 3. To keep matters as simple as possible, all the logarithmic structures we use below are fine and saturated (Definition 3.12). These are the logarithmic structures closest to toroidal embeddings which still allow us to pass to arbitrary subschemes.

The purpose of this text is to survey a number of ways one can think of generalizing fans of toric varieties to the realm of logarithmic structures: Kato fans, Artin fans, polyhedral cone complexes, and skeletons. These are all closely related and, under appropriate assumptions, equivalent. But the different ways they are realized provide for entirely different applications. We do not address skeletons of logarithmic structures over non-trivially valued non-Archimedean fields—see Werner’s contribution [63]. Nor do we address the general theory of skeletons of Berkovich spaces and their generalizations [13, 28, 40].

1.2 Kato Fans

In [34], Kato associated a combinatorial structure F_X , which has since been called *the Kato fan of X* , to a logarithmically regular scheme X (with Zariski-local charts), see Sect. 4. This is a reformulation of the polyhedral cone complexes of [35] which realizes them within the category of monoidal spaces. In [59], one further generalizes the construction of the associated Kato fan to *logarithmic structures without monodromy*. As in [35], Kato fans provide a satisfying theory encoding logarithmically smooth birational modifications in terms of subdivisions of Kato fans. Procedures for resolution of singularities of polyhedral cone complexes of

Kato fans are given in [35] and [34]. As an outcome, we obtain a combinatorial procedure for resolution of singularities of logarithmically smooth structures without self-intersections.

Kato fans, or generalizations of them, can be constructed for more general logarithmic structures. We briefly discuss such a construction using sheaves on the category of Kato fans. A different, and possibly more natural approach, is obtained using Artin fans.

1.3 Artin Fans and Olsson’s Stack of Logarithmic Structures

In order to import notions and structures from scheme theory to logarithmic geometry, Olsson [46] showed that a logarithmic structure X on a given underlying scheme \underline{X} is equivalent to a morphism $\underline{X} \rightarrow \underline{\mathbf{Log}}$, where $\underline{\mathbf{Log}}$ is a rather large, zero-dimensional Artin stack—the moduli stack of logarithmic structures. It carries a universal logarithmic structure whose associated logarithmic algebraic stack we denote by \mathbf{Log} —providing a universal family of logarithmic structures $\mathbf{Log} \rightarrow \underline{\mathbf{Log}}$.

Being universal, the logarithmic stack \mathbf{Log} cannot reflect the combinatorics of X . In Sect. 5 we define, following [3, 8], the notion of *Artin fans*, and show that the morphism $X \rightarrow \mathbf{Log}$ factors through an initial morphism $X \rightarrow \mathcal{A}_X$, where \mathcal{A}_X is an Artin fan and $\mathcal{A}_X \rightarrow \mathbf{Log}$ is étale and representable.

We argue that, unlike \mathbf{Log} , the Artin fan \mathcal{A}_X is a combinatorial object which encodes the combinatorial structure of X . Indeed, when X is without monodromy, the underlying topological space of \mathcal{A}_X is simply the Kato fan F_X . In other words, \mathcal{A}_X combines the advantages of the Kato fan F_X , being combinatorial, and of \mathbf{Log} , being algebraic.

In addition \mathcal{A}_X exists in greater generality, when X is allowed to have self-intersections and monodromy—even not to be logarithmically smooth; in a roundabout way, it provides a definition of F_X in this generality, by taking the underlying “monoidal space”—or more properly, “monoidal stack”—of \mathcal{A}_X .

The theory of Artin fans is not perfect. Its current foundations lack full functoriality of the construction of $X \rightarrow \mathcal{A}_X$, just as Olsson’s characteristic morphism $X \rightarrow \mathbf{Log}$ is functorial only for strict morphisms $Y \rightarrow X$ of logarithmic schemes. In Sect. 5.4.2 we provide a patch for this problem, again following Olsson’s ideas.

1.4 Artin Fans and Unobstructed Deformations

Artin fans were developed in [3, 8] and the forthcoming [6] in order to study logarithmic Gromov–Witten theory. The idea is that since an Artin fan \mathcal{A} is logarithmically étale, a map $f : C \rightarrow \mathcal{A}_X$ from a curve to \mathcal{A}_X is logarithmically unobstructed. Precursors to this result for specific X were obtained in [4, 5, 7, 16]. In [5] an approach to Jun Li’s *expanded degenerations* was provided using what

in hindsight we might call the Artin fan of the affine line $\mathcal{A} = \mathcal{A}_{\mathbb{A}^1}$. The papers [4, 7, 16] use this formalism to prove comparison results in relative Gromov–Witten theory. For logarithmically smooth X , the map $X \rightarrow \mathcal{A}_X$ was used in [3] to prove that *logarithmic* Gromov–Witten invariants are invariant under logarithmic blowings up; in [8], it was used for general X to complete a proof of boundedness of the space of logarithmic stable maps. We review these results, for which both the combinatorial and algebraic features of \mathcal{A}_X are essential, in Sect. 6. They serve as evidence that the algebraic structure of Artin fans is an advantage over the purely combinatorial structure of their associated Kato fans.

1.5 Skeletons and Tropicalization

In Sect. 7 we follow Thuillier [58] and associate to a Zariski-logarithmically smooth scheme X its extended cone complex $\overline{\Sigma}_X$. This is a variant of the cone complex Σ_X of [35], and is related to the Kato fan in an intriguing manner:

$$\overline{\Sigma}_X = F_X(\mathbb{R}_{\geq 0} \sqcup \{\infty\}).$$

The complex $\overline{\Sigma}_X$ is canonically homeomorphic to the skeleton $\mathfrak{S}(X)$ of the non-Archimedean space X^{\natural} associated to the logarithmic scheme X , when viewing the base field as a valued field with trivial valuation, as developed by Thuillier [58]. In this case there is a continuous map $X^{\natural} \rightarrow \overline{\Sigma}_X$, and $\overline{\Sigma}_X \subset X^{\natural}$ is a strong deformation retract. Thuillier used this formalism to prove a compelling result independent of logarithmic or non-Archimedean considerations: the homotopy type of the dual complex of a logarithmic resolution of singularities does not depend on the choice of resolution.

For a general fine and saturated logarithmic scheme X , we still have a continuous mapping $X^{\natural} \rightarrow \overline{\Sigma}_X$, although we do not have a continuous section $\overline{\Sigma}_X \subset X$. It is argued in [59] that the image of $X^{\natural} \rightarrow \overline{\Sigma}_X$ can be viewed as the *tropicalization* of X .

Note that in this discussion we have limited the base field of X to have a trivial absolute value. A truly satisfactory theory must apply to subvarieties defined over valued-field extensions, in particular with nontrivial valuation.

1.6 Analytification of Artin Fans

Artin fans can be tied into the skeleton picture via their analytifications, as we indicate in Sect. 8. We again consider our base field as a trivially valued field. The analytification of the morphism $\phi_X : X \rightarrow \mathcal{A}_X$ is a morphism $\phi_X^{\natural} : X^{\natural} \rightarrow \mathcal{A}_X^{\natural}$ into an analytic Artin stack \mathcal{A}_X^{\natural} . The whole structure sits in a commutative diagram

$$\begin{array}{ccccc}
 X^\natural & \longrightarrow & \mathcal{A}_X^\natural & \longrightarrow & \overline{\Sigma}_X \\
 \rho_X \left(\begin{array}{c} \downarrow \\ \downarrow \end{array} \right) r_X & & \rho_{\mathcal{A}} \left(\begin{array}{c} \downarrow \\ \downarrow \end{array} \right) r_{\mathcal{A}} & & \rho_{\Sigma} \left(\begin{array}{c} \downarrow \\ \downarrow \end{array} \right) r_{\Sigma} \\
 X & \longrightarrow & \mathcal{A}_X & \longrightarrow & F_X
 \end{array}$$

On the left side of the diagram the horizontal arrows remain in their respective categories—algebraic on the bottom, analytic on top—but discard all geometric data of X and X^\natural except the combinatorics of the logarithmic structure. On the right side the arrows $\mathcal{A}_X^\natural \rightarrow \overline{\Sigma}_X$ and $\mathcal{A}_X \rightarrow F_X$ discard the analytic and algebraic data and preserve topological and monoidal structures. In particular $\mathcal{A}_X^\natural \rightarrow \overline{\Sigma}_X$ is a *homeomorphism*, endowing the familiar complex $\overline{\Sigma}_X$ with an analytic stack structure.

1.7 Into the Future

While we have provided evidence that the algebraic structure of \mathcal{A}_X has advantages over the underlying monoidal structure F_X , at this point we can only hope that the analytic structure \mathcal{A}_X^\natural would have significant advantages over the underlying piecewise-linear structure $\overline{\Sigma}_X$, as applications are only starting to emerge, see [55].

We discuss some questions that might usher further applications in Sect. 9.

2 Toric Varieties and Toroidal Embeddings

Mostly for convenience, we work here over an algebraically closed field k . We recall, in briefest terms, the standard setup of toric varieties and toroidal embeddings.


2.1 Toric Varieties

Consider a torus $T \simeq \mathbb{G}_m^n$ and a normal variety X on which T acts with a dense orbit isomorphic to T —and fix such isomorphism. Write M for the character group of T and write N for the co-character group. Then M and N are both isomorphic to \mathbb{Z}^n and are canonically dual to each other. We write $M_{\mathbb{R}} = M \otimes \mathbb{R}$ and $N_{\mathbb{R}} = N \otimes \mathbb{R}$ for the associated real vector spaces.

2.2 Affine Toric Varieties and Cones

If X is affine, it is canonically isomorphic to $X_\sigma = \text{Spec } k[S_\sigma]$, where $\sigma \subset N_{\mathbb{R}}$ is a strictly convex rational polyhedral cone, and $S_\sigma = M \cap \sigma^\vee$ is the monoid of lattice

Fig. 1 The fan of \mathbb{P}^2



points in the n -dimensional cone $\sigma^\vee \subset M_{\mathbb{R}}$ dual to σ . Such affine X_σ contains a unique closed orbit \mathcal{O}_σ , which is itself isomorphic to a suitable quotient torus of T .

2.3 Invariant Opens

A nonempty torus-invariant affine open subset of X_σ is always of the form X_τ where $\tau \prec \sigma$ is a face of σ —either σ itself or the intersection of σ with a supporting hyperplane. For instance, the torus T itself corresponds to the vertex $\{0\}$ of σ .

2.4 Fans

Any toric variety X is covered by invariant affine opens of the form X_{σ_i} , and the intersection of X_{σ_i} with X_{σ_j} is of the form $X_{\tau_{ij}}$ for a common face $\tau_{ij} = \sigma_i \cap \sigma_j$. It follows that the cones σ_i form a *fan* Δ_X in N . This means precisely that Δ_X is a collection of strictly convex rational polyhedral cones in N , that any face of a cone in Δ_X is a member of Δ_X , and the intersection of any two cones in Δ_X is a common face (see Fig. 1).

And vice versa: given a fan Δ in N one can glue together the associated affine toric varieties X_σ along the affine opens X_τ to form a toric variety $X(\Delta)$.

2.5 Categorical Equivalence

One defines a morphism from a toric variety $T_1 \subset X_1$ to another $T_2 \subset X_2$ to be an equivariant morphism $X_1 \rightarrow X_2$ extending a homomorphism of tori $T_1 \rightarrow T_2$. On the other hand one defines a morphism from a fan Δ_1 in $(N_1)_{\mathbb{R}}$ to a fan Δ_2 in $(N_2)_{\mathbb{R}}$ to be a group homomorphism $N_1 \rightarrow N_2$ sending each cone $\sigma \in \Delta_1$ into some cone $\tau \in \Delta_2$.

A fundamental theorem says

Theorem 2.1. *The correspondence above extends to an equivalence of categories between the category of toric varieties over k and the category of fans.*

Under this equivalence, toric birational modifications $X_1 \rightarrow X_2$ of X_2 correspond to subdivisions $\Delta_{X_1} \rightarrow \Delta_{X_2}$ of Δ_{X_2} .


2.6 Extended Fans

The closure $\overline{\mathcal{O}}_\sigma \subset X$ of a T -orbit \mathcal{O}_σ , which is itself a toric variety, is encoded in Δ_X , but in a somewhat cryptic manner. Thuillier [58] provided a way to add all the fans of these loci $\overline{\mathcal{O}}_\sigma$ and obtain a compactification $\Delta_X \subset \overline{\Delta}_X$: instead of gluing together the cones σ along their faces, one replaces σ with a natural compactification, the *extended cone*

$$\overline{\sigma} := \text{Hom}_{\text{Mon}}(S_\sigma, \mathbb{R}_{\geq 0} \cup \{\infty\}),$$

where the notation Hom_{Mon} stands for the set of monoid homomorphisms. This has the effect of adding, in one step, lower dimensional cones isomorphic to $\sigma / \text{Span } \tau$ at infinity corresponding to the closure of \mathcal{O}_τ in X_σ , for all $\tau \prec \sigma$. We still have that $\overline{\tau}_{ij} = \overline{\sigma}_i \cap \overline{\sigma}_j$, and one can glue together these extended cones to obtain the *extended fan* $\overline{\Delta}_X$ (Fig. 2).

Fig. 2 The extended fan of \mathbb{P}^2



2.7 Toroidal Embeddings

The theory of toroidal embedding was developed in [35] in order to describe varieties that look locally like toric varieties. A toroidal embedding $U \subset X$ is a dense open subset of a normal variety X such that, for any closed point x , the completion $\widehat{U}_x \subset \widehat{X}_x$ is isomorphic to the completion of an affine toric variety $T_x \subset X_{\sigma_x}$. Equivalently, each $x \in X$ should admit an étale neighborhood $\phi_x : V_x \rightarrow X$ and an étale morphism $\psi_x : V_x \rightarrow X_{\sigma_x}$ such that $\psi_x^{-1}T_x = \phi_x^{-1}U$. Note that the open set $U \subset X$ serves as a global structure connecting the local pictures $T_x \subset X_{\sigma_x}$.

2.8 The Cone Complex of a Toroidal Embedding Without Self-Intersections


If the morphisms $\phi_x : V_x \rightarrow X$ are assumed to be Zariski open embeddings, then the toroidal embedding is a *toroidal embedding without self-intersections*. In this case the book [35] provides a polyhedral cone complex Σ_X replacing the fan of a toric variety. The main difference is that the cones of the complex Σ_X do not lie linearly inside an ambient space of the form $N_{\mathbb{R}}$.

For a toroidal embedding without self-intersections, the strata of X_{σ_x} glue together to form a stratification $\{\mathcal{O}_i\}$ of X . For $x \in \mathcal{O}_i$ the cone σ_x , along with its sublattice $\sigma_x \cap N$, is independent of x . It can be described canonically as follows. Let X_i be the *star* of \mathcal{O}_i , namely the union of strata containing \mathcal{O}_i in their closures. It is an open subset of X . Let \overline{M}_i be the monoid of effective Cartier divisors on X_i supported on $X_i \setminus U$. Let $N_\sigma = \text{Hom}_{\text{Mon}}(\overline{M}_i, \mathbb{N})$ be the dual monoid. Then $\sigma_x = (N_\sigma)_{\mathbb{R}}$ is the associated cone. When one passes to another stratum contained in X_i one obtains a face $\tau < \sigma$, and $N_\tau = \tau \cap N_\sigma$. These cones glue together naturally, in a manner compatible with the sublattices, to form a cone complex Σ_X with integral structure. Unlike the case of fans, the intersection $\sigma_i \cap \sigma_j$ could be a whole common subfan of σ_i and σ_j , and not necessarily one cone (Fig. 3).

2.9 Extended Complexes

Just as in the case of toric varieties, these complexes canonically admit compactifications $\Sigma_X \subset \overline{\Sigma}_X$, obtained by replacing each cone σ_x by the associated extended cone $\overline{\sigma}_x$. This structure was introduced in Thuillier’s [58].

Fig. 3 The complex of \mathbb{P}^2 with the divisor consisting of a line and a transverse conic: the cones associated to the zero-dimensional strata meet along both their edges



2.10 Functoriality

Let $U_1 \subset X_1$ and $U_2 \subset X_2$ be toroidal embeddings without self-intersections and $f : X_1 \rightarrow X_2$ a dominant morphism such that $f(U_1) \subset U_2$. Then one canonically obtains a mapping $\Sigma_{X_1} \rightarrow \Sigma_{X_2}$, simply because Cartier divisors supported away from U_2 pull back to Cartier divisors supported away from U_1 . This mapping is continuous, sends cones into cones linearly, and sends lattice points to lattice points. Declaring such mappings to be mappings of polyhedral cone complexes with integral structure, we obtain a functor from toroidal embeddings to polyhedral cone complexes. This functor is far from being an equivalence.

In [35] one focuses on *toroidal modifications* $f : X_1 \rightarrow X_2$, namely those birational modifications described on charts of X_2 by toric modifications of the toric varieties X_{σ_x} . Then one shows

Theorem 2.2. *The correspondence $X_1 \mapsto \Sigma_{X_1}$ extends to an equivalence of categories between toroidal modifications of $U_2 \subset X_2$, and subdivisions of Σ_{X_2} .*

3 Logarithmic Structures

We briefly review the theory of logarithmic structures [33].

3.1 Notation for Monoids

Definition 3.1. A *monoid* is a commutative semi-group with a unit. A morphism of monoids is required to preserve the unit element.

We denote the category of monoids by the symbol **Mon**.

Given a monoid P , we can associate a group

$$P^{gp} := \{(a, b) \mid (a, b) \sim (c, d) \text{ if } \exists s \in P \text{ such that } s + a + d = s + b + c\}.$$

Note that any morphism from P to an abelian group factors through P^{gp} uniquely.

Definition 3.2. A monoid P is called *integral* if $P \rightarrow P^{gp}$ is injective. It is called *fine* if it is integral and finitely generated.

An integral monoid P is said to be *saturated* if whenever $p \in P^{gp}$ and n is a positive integer such that $np \in P$ then $p \in P$.

As has become customary, we abbreviate the combined condition “fine and saturated” to *fs*.

3.2 Logarithmic Structures

Definition 3.3. Let \underline{X} be a scheme. A *pre-logarithmic structure* on \underline{X} is a sheaf of monoids M_X on the small étale site $\text{ét}(\underline{X})$ combined with a morphism of sheaves of monoids: $\alpha : M_X \rightarrow \mathcal{O}_{\underline{X}}$, called the *structure morphism*, where we view $\mathcal{O}_{\underline{X}}$ as a monoid under multiplication. A pre-logarithmic structure is called a *logarithmic structure* if $\alpha^{-1}(\mathcal{O}_{\underline{X}}^*) \cong \mathcal{O}_{\underline{X}}^*$ via α . The pair (\underline{X}, M_X) is called a *logarithmic scheme*, and will be denoted by X .

The structure morphism α is frequently denoted exp and an inverse $\mathcal{O}_X^* \rightarrow M_X$ is denoted log .

Definition 3.4. Given a logarithmic scheme X , the quotient sheaf $\overline{M}_X = M_X/\mathcal{O}_X^*$ is called the *characteristic monoid*, or just the *characteristic*, of the logarithmic structure M_X .

Definition 3.5. Let M and N be pre-logarithmic structures on \underline{X} . A *morphism* between them is a morphism $M \rightarrow N$ of sheaves of monoids which is compatible with the structure morphisms.

Definition 3.6. Let $\alpha : M \rightarrow \mathcal{O}_{\underline{X}}$ be a pre-logarithmic structure on \underline{X} . We define the *associated logarithmic structure* M^a to be the push-out of

$$\begin{array}{ccc} \alpha^{-1}(\mathcal{O}_{\underline{X}}^*) & \longrightarrow & M \\ \downarrow & & \\ \mathcal{O}_{\underline{X}}^* & & \end{array}$$

in the category of sheaves of monoids on $\text{ét}(\underline{X})$, endowed with

$$M^a \rightarrow \mathcal{O}_{\underline{X}} \quad (a, b) \mapsto \alpha(a)b \quad (a \in M, b \in \mathcal{O}_{\underline{X}}^*).$$

The following are two standard examples from [33, (1.5)]:

Example 3.7. Let \underline{X} be a smooth scheme with an effective divisor $D \subset \underline{X}$. Then we have a standard logarithmic structure M on \underline{X} associated to the pair (\underline{X}, D) , where

$$M_X := \{f \in \mathcal{O}_{\underline{X}} \mid f|_{\underline{X} \setminus D} \in \mathcal{O}_{\underline{X}}^*\}$$

with the structure morphism $M_X \rightarrow \mathcal{O}_{\underline{X}}$ given by the canonical inclusion. This is already a logarithmic structure, as any section of $\mathcal{O}_{\underline{X}}^*$ is already in M_X .

Example 3.8. Let P be an fs monoid, and $\underline{X} = \text{Spec } \mathbb{Z}[P]$ be the associated affine toric scheme. Then we have a standard logarithmic structure M_X on \underline{X} associated to the pre-logarithmic structure

$$P \rightarrow \mathbb{Z}[P]$$

defined by the obvious inclusion.

We denote by $\text{Spec}(P \rightarrow \mathbb{Z}[P])$ the log scheme (\underline{X}, M_X) .

3.3 Inverse Images

Let $f : \underline{X} \rightarrow \underline{Y}$ be a morphism of schemes. Given a logarithmic structure M_Y on \underline{Y} , we can define a logarithmic structure on \underline{X} , called the *inverse image* of M_Y , to be the logarithmic structure associated to the pre-logarithmic structure $f^{-1}(M_Y) \rightarrow f^{-1}(\mathcal{O}_Y) \rightarrow \mathcal{O}_X$. This is usually denoted by $f^*(M_Y)$. Using the inverse image of logarithmic structures, we can give the following definition.

Definition 3.9. A morphism of logarithmic schemes $X \rightarrow Y$ consists of a morphism of underlying schemes $f : \underline{X} \rightarrow \underline{Y}$, and a morphism $f^b : f^*M_Y \rightarrow M_X$ of logarithmic structures on \underline{X} . The morphism is said to be *strict* if f^b is an isomorphism.

We denote by **LogSch** the category of logarithmic schemes.

3.4 Charts of Logarithmic Structures

Definition 3.10. Let X be a logarithmic scheme, and P a monoid. A *chart* for M_X is a morphism $P \rightarrow \Gamma(X, M_X)$, such that the induced map of logarithmic structures $P^a \rightarrow M_X$ is an isomorphism, where P^a is the logarithmic structure associated to the pre-logarithmic structure given by $P \rightarrow \Gamma(X, M_X) \rightarrow \Gamma(X, \mathcal{O}_X)$.

In fact, a chart of M_X is equivalent to a morphism

$$f : X \rightarrow \text{Spec}(P \rightarrow \mathbb{Z}[P])$$

such that f^b is an isomorphism. In general, we have the following:

Lemma 3.11 ([45, 1.1.9]). *The mapping*

$$\text{Hom}_{\text{LogSch}}(X, \text{Spec}(P \rightarrow \mathbb{Z}[P])) \rightarrow \text{Hom}_{\text{Mon}}(P, \Gamma(X, M_X))$$

associating to f the composition

$$P \rightarrow \Gamma(\underline{X}, P_X) \xrightarrow{\Gamma(f^b)} \Gamma(\underline{X}, M_X)$$

is a bijection.

We will see in Example 5.2 that Artin cones have a similar universal property on the level of characteristic monoids.

Definition 3.12. A logarithmic scheme X is said to be *fine*, if étale locally there is a chart $P \rightarrow M_X$ with P a fine monoid. If moreover P can be chosen to be saturated, then X is called a *fine and saturated* (or *fs*) logarithmic structure. Finally, if P can be chosen isomorphic to \mathbb{N}^k , we say that the logarithmic structure is *locally free*.

Lemma 3.13. *Let X be a fine and saturated logarithmic scheme. Then for any geometric point $\bar{x} \in X$, there exists an étale neighborhood $U \rightarrow X$ of \bar{x} with a chart $\bar{M}_{\bar{x},X} \rightarrow M_U$, such that the composition $\bar{M}_{\bar{x},X} \rightarrow M_U \rightarrow \bar{M}_{\bar{x},X}$ is the identity.*

Proof. This is a special case of [46, Proposition 2.1]. □

3.5 Logarithmic Differentials

To form sheaves of logarithmic differentials, we add to the sheaf $\Omega_{\underline{X}/\underline{Y}}$ symbols of the form $d \log(\alpha(m))$ for all elements $m \in M_X$, as follows:

Definition 3.14. Let $f : X \rightarrow Y$ be a morphism of fine logarithmic schemes. We introduce the sheaf of relative logarithmic differentials $\Omega^1_{X/Y}$ given by

$$\Omega^1_{X/Y} = \left(\Omega_{\underline{X}/\underline{Y}} \oplus (\mathcal{O}_{\underline{X}} \otimes_{\mathbb{Z}} M_X^{gp}) \right) / \mathcal{K}$$

where \mathcal{K} is the $\mathcal{O}_{\underline{X}}$ -module generated by local sections of the following forms:

- (1) $(d\alpha(a), 0) - (0, \alpha(a) \otimes a)$ with $a \in M_X$ and
- (2) $(0, 1 \otimes a)$ with $a \in \text{im}(f^{-1}(M_Y) \rightarrow M_X)$.

The universal derivation (∂, D) is given by $\partial : \mathcal{O}_{\underline{X}} \xrightarrow{d} \Omega_{\underline{X}/\underline{Y}} \rightarrow \Omega^1_{X/Y}$ and $D : M_X \rightarrow \mathcal{O}_{\underline{X}} \otimes_{\mathbb{Z}} M_X^{gp} \rightarrow \Omega^1_{X/Y}$.

Example 3.15. Let $h : Q \rightarrow P$ be a morphism of fine monoids. Denote $X = \text{Spec}(P \rightarrow \mathbb{Z}[P])$ and $Y = \text{Spec}(Q \rightarrow \mathbb{Z}[Q])$. Then we have a morphism $f : X \rightarrow Y$ induced by h . A direct calculation shows that $\Omega^1_f = \mathcal{O}_{\underline{X}} \otimes_{\mathbb{Z}} \text{coker}(h^{gp})$. The free generators correspond to the logarithmic differentials $d \log(\alpha(p))$ for $p \in P$, which are regular on the torus $\text{Spec } \mathbb{Z}[P^{gp}]$, modulo those coming from Q . This can also be seen from the universal property of the sheaf of logarithmic differentials.

3.6 Logarithmic Smoothness

Consider the following commutative diagram of logarithmic schemes illustrated with solid arrows:

$$\begin{array}{ccc}
 T_0 & \xrightarrow{\phi} & X \\
 J \downarrow j & \nearrow & \downarrow f \\
 T_1 & \xrightarrow{\psi} & Y
 \end{array}
 \tag{1}$$

where j is a *strict* closed immersion (Definition 3.9) defined by the ideal J with $J^2 = 0$. We define logarithmic smoothness by the infinitesimal lifting property:

Definition 3.16. A morphism $f : X \rightarrow Y$ of fine logarithmic schemes is called *logarithmically smooth* (resp., *logarithmically étale*) if the underlying morphism $\underline{X} \rightarrow \underline{Y}$ is locally of finite presentation and for any commutative diagram (1), étale locally on T_1 there exists a (resp., there exists a unique) morphism $g : T_1 \rightarrow X$ such that $\phi = g \circ j$ and $\psi = f \circ g$.

We have the following useful criterion for smoothness from [33, Theorem 3.5].

Theorem 3.17 (K. Kato). *Let $f : X \rightarrow Y$ be a morphism of fine logarithmic schemes. Assume we have a chart $Q \rightarrow M_Y$, where Q is a finitely generated integral monoid. Then the following are equivalent:*

- (1) f is logarithmically smooth (resp., logarithmically étale) and
- (2) étale locally on X , there exists a chart $(P_X \rightarrow M_X, Q_Y \rightarrow M_Y, Q \rightarrow P)$ extending the chart $Q_Y \rightarrow M_Y$, satisfying the following properties:
 - (a) The kernel and the torsion part of the cokernel (resp., the kernel and the cokernel) of $Q^{gp} \rightarrow P^{gp}$ are finite groups of orders invertible on X .
 - (b) The induced morphism from $\underline{X} \rightarrow \underline{Y} \times_{\text{Spec } \mathbb{Z}[Q]} \text{Spec } \mathbb{Z}[P]$ is étale in the usual sense.

Remark 3.18. (1) We can require $Q^{gp} \rightarrow P^{gp}$ in (a) to be injective, and replace the requirement that $\underline{X} \rightarrow \underline{Y} \times_{\text{Spec } \mathbb{Z}[Q]} \text{Spec } \mathbb{Z}[P]$ be étale in (b) by requiring it to be smooth without changing the conclusion of Theorem 3.17.

(2) In this theorem something wonderful happens, which Kato calls “the magic of log.” The arrow in (b) shows that a logarithmically smooth morphism is “locally toric” relative to the base. Consider the case where Y is a logarithmic scheme

with underlying space given by $\text{Spec } \mathbb{C}$ with the trivial logarithmic structure, and $X = \text{Spec}(P \rightarrow \mathbb{C}[P])$ where P is a fine, saturated, and torsion free monoid. Then \underline{X} is a toric variety with the action of $\text{Spec } \mathbb{C}[P^{gp}]$. According to the theorem, X is logarithmically smooth relative to Y , though the underlying space might be singular. These singularities are called toric singularities in [34]. This is closely related to the classical notion of toroidal embeddings [35].

Logarithmic differentials behave somewhat analogously to differentials:

Proposition 3.19. *Let $X \xrightarrow{f} Y \xrightarrow{g} Z$ be a sequence of morphisms of fine logarithmic schemes.*

- (1) *There is a natural exact sequence $f^* \Omega_g^1 \rightarrow \Omega_{fg}^1 \rightarrow \Omega_f^1 \rightarrow 0$.*
- (2) *If f is logarithmically smooth, then Ω_f^1 is a locally free \mathcal{O}_X -module, and we have the following exact sequence: $0 \rightarrow f^* \Omega_g^1 \rightarrow \Omega_{fg}^1 \rightarrow \Omega_f^1 \rightarrow 0$.*
- (3) *If gf is logarithmically smooth and the sequence in (2) is exact and splits locally, then f is logarithmically smooth.*

A proof can be found in [45, Chap. IV].

4 Kato Fans and Resolution of Singularities

4.1 The Monoidal Analogues of Schemes


In parallel to the theory of schemes, Kato developed a theory of fans, with the role of commutative rings played by monoids. As with schemes, the theory begins with the spectrum of a monoid:

Definition 4.1 ([34, Definition (5.1)]). Let M be a monoid. A subset $I \subset M$ is called an *ideal* of M if $M + I \subset I$. If $M \setminus I$ is a submonoid of M , then I is called a *prime ideal* of M . The set of prime ideals of M is denoted $\text{Spec } M$ and called the *spectrum* of M .

If $f : M \rightarrow N$ is a homomorphism of monoids and $P \subset N$ is a prime ideal, then $f^{-1}P \subset M$ is a prime ideal as well. Therefore f induces a morphism of spectra: $\text{Spec } N \rightarrow \text{Spec } M$.

Definition 4.2 ([34, Definition (5.2)]). Suppose that M is a monoid and S is a subset of M . We write $M[-S]$ for the initial object among the monoids N equipped a morphism $f : M \rightarrow N$ such that $f(S)$ is invertible. When S consists of a single element s , we also write $M[-s]$ in lieu of $M[-\{s\}]$.

Fig. 4 Points and topology of $\text{Spec } \mathbb{N}^2$. The big point corresponds to the ideal \emptyset , the intermediate points are the primes $\mathbb{N} \times \mathbb{N}_{>0}$ and $\mathbb{N}_{>0} \times \mathbb{N}$, and the small closed point is the maximal ideal $\mathbb{N}_{>0}^2$. The arrows indicate specialization, thus determining the topology



It is not difficult to construct $M[-S]$ with the familiar Grothendieck group construction $M \mapsto M^{gp}$ of Definition 3.1. Certainly $M[-S]$ coincides with $M[-S']$ where S' is the submonoid of M generated by S . One may therefore assume that S is a submonoid of M . Then for $M[-S]$ one may take the set of formal differences $m - s$ with $m \in M$ and $s \in S$, subject to the familiar equivalence relation:

$$m - s \sim m' - s' \iff \exists t \in S, t + m + s' = t + m' + s$$

If M is integral, then one may construct $M[-S]$ as a submonoid of M^{gp} .

The topology of the spectrum of a monoid is defined exactly as for schemes:

Definition 4.3 ([34, Definition (9.2)]). Let M be a monoid. For any $f \in M$, let $D(f) \subset \text{Spec } M$ be the set of prime ideals $P \subset M$ such that $f \notin P$. A subset of $\text{Spec } M$ is called open if it is open in the minimal topology in which the $D(f)$ are open subsets.

Equivalently $D(f)$ is the image of the map $\text{Spec}(M[-f]) \rightarrow \text{Spec } M$. The intersection of $D(f)$ and $D(g)$ is $D(f + g)$ so the sets $D(f)$ form a basis for the topology of $\text{Spec } M$ (Fig. 4).

We equip $\text{Spec } M$ with a sheaf of monoids $\mathcal{M}_{\text{Spec } M}$ where

$$\mathcal{M}_{\text{Spec } M}(D(f)) = M[-f]/M[-f]^*$$

where $M[-f]^*$ is the set of invertible elements of M . This is a sharply monoidal space:

Definition 4.4 ([34, Definition (9.1)]). Recall that a monoid is called *sharp* if its only invertible element is the identity element 0. If M and N are sharp monoids, then a sharp homomorphism $f : M \rightarrow N$ is a homomorphism of monoids such that $f^{-1}\{0\} = \{0\}$.

A *sharply monoidal space* is a pair (S, \mathcal{M}_S) , where S is a topological space and \mathcal{M}_S is a sheaf of sharp monoids on S . A morphism of sharply monoidal spaces $f : (S, \mathcal{M}_S) \rightarrow (T, \mathcal{M}_T)$ consists of a continuous function $f : S \rightarrow T$ and a sharp homomorphism of sheaves of sharp monoids $f^{-1}\mathcal{M}_T \rightarrow \mathcal{M}_S$.

A sharply monoidal space is called an *affine Kato fan* or a *Kato cone* if it is isomorphic to $(\text{Spec } M, \mathcal{M}_{\text{Spec } M})$. A sharply monoidal space is called a *Kato fan* if it admits an open cover by Kato cones.

We call a Kato fan *integral* or *saturated* if it admits a cover by the spectra of monoids with the respective properties (Definition 3.2). A Kato fan is called *locally of finite type* if it admits a cover by spectra of finitely generated monoids. We use *fine* as a synonym for integral and locally of finite type.

A large collection of examples of Kato fans is obtained from fans of toric varieties (see Sect. 4.3) or logarithmically smooth schemes (Sect. 4.5). In particular, any toric singularity is manifested in a Kato fan.

Any Kato cone $\text{Spec } M$ of a finitely generated integral monoid M contains an open point corresponding to the ideal $\emptyset \subset M$, carrying the trivial stalk $M^{\text{gp}}/M^{\text{gp}} = 0$. We can always glue an arbitrary collection of such Kato cones along their open points. If the collection is infinite, this gives examples of connected Kato fans which are not quasi-compact.

4.2 Points and Kato Cones

Definition 4.5. Let $F' \rightarrow F$ be a morphism of fine, saturated Kato fans. The morphism is said to be *quasi-compact* if the preimage of any open subcone of F is quasi-compact.

Lemma 4.6. *Let F be a Kato fan. There is a bijection between the open Kato subcones of F and the points of F .*

Proof. Suppose that $U = \text{Spec } M$ is an open subcone of F . Let $P \subset M$ be the complement of $0 \in M$. Then P is a prime ideal, hence corresponds to a point of F .

To give the inverse, we show that every point of F has a minimal open affine neighborhood. Indeed, suppose that U is an open affine neighborhood of $p \in F$. Then the underlying topological space of U is finite, so there is a smallest open subset of U containing p . Replace U with this open subset. It must be affine, since affine open subsets form a basis for the topology of U .

It is straightforward to see that these constructions are inverse to one another. \square

Lemma 4.7. *A morphism of fine, saturated Kato fans is quasi-compact if and only if it has finite fibers. In particular, a Kato fan is quasi-compact if and only if its underlying set is finite.*

4.3 From Fans to Kato Fans

If Δ is a fan in $N_{\mathbb{R}}$ in the sense of toric geometry, it gives rise to a Kato fan F_{Δ} . The underlying topological space of F_{Δ} is the set of cones of Δ and a subset is open if and only if it contains all the faces of its elements. In particular, if σ is one

of the cones of Δ , then the set of faces of σ is an open subset F_σ of K and these open subsets form a basis. We set $\mathcal{M}(F_\sigma) = (M \cap \sigma^\vee)/(M \cap \sigma^\vee)^*$ where M is the dual lattice of N and σ^\vee is the dual cone of σ . With this sheaf of monoids, $F_\sigma \simeq \text{Spec}(M \cap \sigma^\vee)/(M \cap \sigma^\vee)^*$, so F_Δ has an open cover by Kato cones, hence is a Kato fan.

4.4 Resolution of Singularities of Kato Fans

Kato introduced the following monoidal space analogue of subdivisions of fans of toric varieties, in such a way that a subdivision of a fan Σ gives rise to a subdivision of the Kato fan F_Σ . A morphism of fans $\Sigma_1 \rightarrow \Sigma_2$ is a subdivision if and only if it induces a bijection on the set of lattice points $\cup_\sigma N_\sigma$. The Kato fan notion is the direct analogue:

Definition 4.8. A morphism $p : F' \rightarrow F$ of fine, saturated Kato fans is called a *proper subdivision* if it is quasi-compact and the morphism

$$\text{Hom}(\text{Spec } \mathbb{N}, F') \rightarrow \text{Hom}(\text{Spec } \mathbb{N}, F)$$

is a bijection.

Remark 4.9. This definition has an appealing resemblance to the valuative criterion for properness.

Explicitly subdividing Kato fans is necessarily less intuitive than subdividing fans. The following examples, which are in direct analogy to subdivisions of fans, may help in developing intuition:

Example 4.10. (i) Suppose that F is a fine, saturated Kato fan and $v : \text{Spec } \mathbb{N} \rightarrow F$ is a morphism. The *star subdivision* of F along v is constructed as follows: If $U = \text{Spec } M$ is an open Kato subcone of F not containing v , then U is included as an open Kato subcone of F' ; if U does contain v , then for each face $V = \text{Spec } N$ contained in U that does not contain v , we include the face $v + V = \text{Spec } M_{v+V}$, where

$$M_{v+V} = \{ \alpha \in M^{\text{gp}} : \alpha|_V \in N \text{ and } v^* \alpha \in \mathbb{N} \}.$$

(ii) Let $F = \text{Spec } M$ be a fine, saturated Kato cone. There is a canonical morphism $\text{Spec } \mathbb{N} \rightarrow F$ by regarding $\text{Hom}(\text{Spec } \mathbb{N}, F) = \text{Hom}(M, \mathbb{N})$ as a monoid and taking the sum of the generators of the 1-dimensional faces of F . This morphism is called the *barycenter* of F .

If F is a fine, saturated Kato fan, we obtain a subdivision $F' \rightarrow F$ by performing star subdivision of F along the barycenters of its open subcones, in decreasing order of dimension. A priori this is well defined if cones of F have bounded dimension, but as the procedure is compatible with restriction to subfans, this works for arbitrary F . This subdivision is called the *barycentric subdivision*.

Definition 4.11. A fine, saturated Kato fan F is said to be *smooth* if it has an open cover by Kato cones $U \simeq \text{Spec } \mathbb{N}^r$.

Theorem 4.12 ([35, Theorem I.11], [23, Sect. 2.6], [34, Proposition (9.8)]). *Let F be a fine, saturated Kato fan. Then there is a proper subdivision $F' \rightarrow F$ such that F' is smooth.*

The classical combinatorial proofs start by first using star or barycentric subdivisions to make the fan simplicial, and then repeatedly reducing the index by further star subdivisions.

4.5 Logarithmic Regularity and Associated Kato Fans

In this section we will work only with logarithmic structures admitting charts *Zariski locally*.

Let X be a logarithmic scheme and x a schematic point of X . Let $I(x, M_X)$ be the ideal of the local ring $\mathcal{O}_{X,x}$ generated by the maximal prime ideal of $M_{X,x}$.

Definition 4.13 ([34, Definition (2.1)]). A locally noetherian logarithmic scheme X admitting Zariski-local charts is called *logarithmically regular* at a schematic point x if it satisfies the following two conditions:

- (i) the local ring $\mathcal{O}_{X,x}/I(x, M_X)$ is regular and
- (ii) $\dim \mathcal{O}_{X,x} = \dim \mathcal{O}_{X,x}/I(x, M) + \text{rank } \overline{M}_{X,x}^{\text{gp}}$.

Example 4.14 ([34, Example (2.2)]). A toric variety with its toric logarithmic structure is logarithmically regular.

If X is a logarithmic scheme, then the Zariski topological space of X is equipped with a sheaf of sharp monoids \overline{M}_X . Thus (X, \overline{M}_X) is a sharply monoidal space. Moreover, morphisms of logarithmic schemes induce morphisms of sharply monoidal spaces. We therefore obtain a functor from the category of logarithmic schemes to the category of sharply monoidal spaces. We may speak in particular about morphisms from logarithmic schemes to Kato fans.

Theorem 4.15 (cf. [34, (10.2)]). *Let X be a fine, saturated locally noetherian, logarithmically regular logarithmic scheme that admits charts Zariski locally. Then there is an initial strict morphism $X \rightarrow F_X$ to a Kato fan.*

We call the Kato fan F_X the *Kato fan associated to X* .

Kato constructs the Kato fan F_X as a skeleton of the logarithmic strata of X . Let $\underline{F}_X \subset X$ be the set of points $x \in X$ such that $I(x, M_X)$ coincides with the maximal ideal of $\mathcal{O}_{X,x}$. As we indicate below these are the generic points of the logarithmic strata of X . Let M_{F_X} be the restriction of \overline{M}_X to \underline{F}_X . We denote the resulting monoidal space by F_X .

Lemma 4.16 ([34, Proposition (10.1)]). *If X is a fine, saturated, locally noetherian, logarithmically regular scheme admitting a chart Zariski locally, then the sharply monoidal space constructed above is a fine, saturated Kato fan.*

Lemma 4.17 ([34, (10.2)]). *There is a canonical continuous retraction of X onto its Kato fan F_X .*

We briefly summarize the definition of the map and omit the rest of the proof. Let x be a point of X . The quotient $\mathcal{O}_{X,x}/I(x, M_X)$ is regular, hence in particular a domain, so it has a unique minimal prime corresponding to a point y of $F_X \subset X$. We define $\pi(x) = y$.

Lemma 4.18. *The Kato fan of X is the initial Kato fan admitting a morphism from X .*

Proof. Let $\varphi : X \rightarrow F$ be another morphism to a Kato fan. By composition with the inclusion $F_X \subset X$ we get a map $F_X \rightarrow F$. We need to show that φ factors through $\pi : X \rightarrow F_X$.

Let x be a point of X and let A be the local ring of x in X . Let $P = \overline{M}_{X,x}$. Let $y = \pi(x)$. Then y is the generic point of the vanishing locus of $I(x, M_X)$. We would like to show $\varphi(y) = \varphi(x)$. We can replace X with $\text{Spec } A$, replace F_X with $F_X \cap \text{Spec } A$, and replace F with an open Kato cone in F containing $\varphi(x)$. Then $F = \text{Spec } Q$ for some fine, saturated, sharp monoid Q and we get a map $Q \rightarrow P$. Since $X \rightarrow F$ is a morphism of sharply monoidal spaces, $X \rightarrow F$ factors through the open subset defined by the kernel of $Q \rightarrow P$, so we can replace F by this open subset and assume that $Q \rightarrow P$ is sharp. But then $X \rightarrow F$ factors through no smaller open subset, so $\varphi(x)$ is the closed point of F . Moreover, we have $\overline{M}_{X,y} = \overline{M}_{X,x}$ so the same reasoning applies to y and shows that $\varphi(y) = \varphi(x)$, as desired. \square

This completes the proof of Theorem 4.15.


4.6 Towards the Monoidal Analogues of Algebraic Spaces

4.6.1 The Need for a More General Approach: The Nodal Cubic

Not every logarithmically smooth scheme has a Kato fan. For example, the divisorial logarithmic structure on \mathbb{A}^2 associated to a nodal cubic curve does not have a chart in any Zariski neighborhood of the node of the cubic. The Kato fan of this logarithmic structure wants to be the Kato fan of the plane with the open subsets corresponding to the complements of the axes glued together (Fig. 5).

This cannot be a Kato fan, because there is no open neighborhood of the closed point that is a Kato cone. Indeed, the closed point has two different generizations to the codimension 1 point.

Fig. 5 The fan of the nodal cubic, drawn as a cone with the two edges glued to each other



Clearly, the solution here is to allow more general types of gluing (colimits) into the definition of a Kato fan. The purpose of this section is to outline what this might entail. Our favorite solution will only come in the next section, where we discuss Artin fans.

The theory of algebraic spaces provides a blueprint for how to proceed. The universal way to add colimits to a category is to pass to its category of presheaves. In order to retain the Kato fans as colimits of their open Kato subcones, we look instead at the category of sheaves. Finally, we restrict attention to those sheaves that resemble Kato cones *étale locally*.

4.6.2 The Need for a More General Approach: The Whitney Umbrella

These *cone spaces* are general enough to include a fan for the divisorial logarithmic structure of the nodal cubic. However, one cannot obtain a fan for the punctured Whitney umbrella this way:


Working over \mathbb{C} , let $X = \mathbb{A}^2 \times \mathbb{G}_m$ with the logarithmic structure pulled back from the standard logarithmic structure on \mathbb{A}^2 , associated to the divisor $xy = 0$ as in Example 3.7. Let $G = \mathbb{Z}/2\mathbb{Z}$ act by $t.(x, y, z) = (y, x, -z)$. Since the divisor is G stable, the action lifts to the logarithmic structure, so the logarithmic structure descends to the quotient $Y = X/G$.

There is a projection $\pi : Y \rightarrow \mathbb{G}_m$, and traversing the nontrivial loop in \mathbb{G}_m results in the automorphism of \mathbb{A}^2 exchanging the axes. Thus the logarithmic structure of Y has *monodromy*.

If there is a map from Y to a cone space Z , then it is not possible for \overline{M}_Y to be pulled back from M_Z . Just as in a Kato fan, the strata of Z on which M_Z is locally constant are discrete and therefore cannot have monodromy.

However, *Deligne–Mumford stacks* can have monodromy at points. By enlarging our perspective to include *cone stacks*, the analogues of Deligne–Mumford stacks for Kato fans, we are able to construct fans that record the combinatorics of the logarithmic strata for any locally connected logarithmic scheme (logarithmic smoothness is not required). The construction proceeds circuitously, by showing that cone stacks form a full subcategory of Artin fans (defined in Sect. 5) and then constructing an Artin fan associated to any logarithmically smooth scheme (Fig. 6).

Fig. 6 The fan of the Whitney umbrella drawn as a cone—the first quadrant—folded over itself via the involution $(x, y) \mapsto (y, x)$



5 Artin Fans

In order to simplify the discussion we work over an algebraically closed field \mathbf{k} .

5.1 Definition and Basic Properties

Definition 5.1. An *Artin fan* is a logarithmic algebraic stack that is logarithmically étale over $\text{Spec } \mathbf{k}$. A morphism of Artin fans is a morphism of logarithmic algebraic stacks. We denote the 2-category of Artin fans by the symbol \mathbf{AF} .

Olsson showed that there is an algebraic stack $\underline{\text{Log}}$ over the category of schemes such that morphisms $\underline{S} \rightarrow \underline{\text{Log}}$ correspond to logarithmic structures S on \underline{S} [46, Theorem 1.1].¹ Equipping $\underline{\text{Log}}$ with its universal logarithmic structure yields the logarithmic algebraic stack $\overline{\text{Log}}$. If \mathcal{X} is an Artin fan, then the morphism $\underline{\mathcal{X}} \rightarrow \underline{\text{Log}}$ induced by the logarithmic structure of \mathcal{X} is étale, and conversely: any logarithmic algebraic stack whose structural morphism to $\overline{\text{Log}}$ is étale is an Artin fan.

Example 5.2. Let V be a toric variety with dense torus T . The toric logarithmic structure of V is T -equivariant, hence descends to a logarithmic structure on $[V/T]$, making $[V/T]$ into an Artin fan.

If $V = \text{Spec } \mathbf{k}[M]$ for some fine, saturated, sharp monoid M , then $[V/T]$ represents the following functor on logarithmic schemes [46, Proposition 5.17]:


$$(X, M_X) \mapsto \text{Hom}(M, \Gamma(X, \overline{M}_X)).$$

As \overline{M}_X is an étale sheaf over X , it is immediate that $[V/T]$ is logarithmically étale over a point.

Definition 5.3. An *Artin cone* is an Artin fan isomorphic to $\mathcal{A}_M = [V/T]$, where $V = \text{Spec } \mathbf{k}[M]$ and M is a fine, saturated, sharp monoid (Fig. 7).

¹Our conventions differ from Olsson’s: in [46], the stack Log parameterizes all fine logarithmic structures; to conform with our convention that all logarithmic structures are fine and saturated, we use the symbol $\overline{\text{Log}}$ to refer to the open substack of Olsson’s stack parameterizing fine and saturated logarithmic structures. This was denoted $\mathcal{T}or$ in [46].

Fig. 7 $\mathcal{A}_{\mathbb{N}^2}$. The small closed point is $B\mathbb{G}_m^2$, the intermediate points are $B\mathbb{G}_m$, and the big point is just a point



If $\mathcal{X} \rightarrow \mathbf{Log}$ is étale and representable, then \mathcal{X} is determined by its étale stack of sections over \mathbf{Log} . Moreover, any étale stack on \mathbf{Log} corresponds to an Artin fan by passage to the espace (champ) étalé. The following lemma characterizes the étale site of \mathbf{Log} as a category of presheaves:

- Lemma 5.4.** (i) *The Artin cones are an étale cover of \mathbf{Log} .*
 (ii) *An Artin cone has no nontrivial representable étale covers.*
 (iii) *If M and N are fine, saturated, sharp monoids, then*

$$\mathrm{Hom}_{\mathbf{AF}}(\mathcal{A}_M, \mathcal{A}_N) = \mathrm{Hom}_{\mathbf{Mon}}(N, M).$$

Proof. Statement (i) was proved in [46, Corollary 5.25]. It follows from the fact that every logarithmic scheme admits a chart étale locally.

Statement (ii) follows from [3, Corollary 2.4.3]. Concretely, étale covers of Artin cones correspond to equivariant étale covers of toric varieties, which restrict to equivariant étale covers of tori, of which there are none other than the trivial ones [3, Proposition 2.4.1].

Statement (iii) follows from $\Gamma(\mathcal{A}_M, \overline{M}_{\mathcal{A}_M}) = M$ and consideration of the functor represented by \mathcal{A}_N (Example 5.2). □

5.2 Categorical Context

Lemma 5.4 (iii) enables us to relate the 2-category of Artin fans to the notions surrounding Kato fans. Let us write **RPC** for the category of rational polyhedral cones. The category **RPC** is equivalent to the opposite of the category of fine, saturated, sharp monoids.² Therefore **RPC** is equivalent to the category of Kato cones and, by Lemma 5.4 (iii), to the category of Artin cones. Furthermore, we obtain

²In fact, the category of fine, saturated, sharp monoids is equivalent to its own opposite. However, we find it helpful to maintain a distinction between cones (which we view as spaces) and monoids (which we view as functions).

Corollary 5.5. *The 2-category of Artin fans with faithful monodromy is fully faithfully embedded in the 2-category of fibered categories over **RPC**.*

Proof. Lemma 5.4 gives an embedding of **AF** in the 2-category of fibered categories on the category of Artin cones and identifies the category of Artin cones with **RPC**. □

This enables us to relate the 2-category **AF** with the framework proposed in Sects. 4.6.1 and 4.6.2: the 2-category of cone stacks suggested in Sect. 4.6.2 is necessarily equivalent to a full subcategory the 2-category **AF**. The key to proving this is the fact that the diagonal of an Artin fan is represented by algebraic spaces, which enables one to relate it to a morphism of “cone spaces” as suggested in Sect. 4.6.1.

In particular, we have a fully faithful embedding **KF** \rightarrow **AF** of the category of Kato fans in the 2-category of Artin fans.

5.3 The Artin Fan of a Logarithmic Scheme

Definition 5.6. A Zariski logarithmic scheme X is said to be *small* with respect to a point $x \in X$, if the restriction morphism $\Gamma(X, \overline{M}) \rightarrow \overline{M}_{X,x}$ is an isomorphism and the closed logarithmic stratum

$$\{y \in X \mid \overline{M}_{X,y} \simeq \overline{M}_{X,x}\}$$

is *connected*. We say X is *small* if it is small with respect to some point.

Let $N = \Gamma(X, \overline{M}_X)$. There is a canonical morphism

$$X \rightarrow \mathcal{A}_N$$

corresponding to the morphism $N \rightarrow \Gamma(X, \overline{M}_X)$. It is shown in [8] that if X is *small* this morphism is initial among all morphisms from X to Artin fans. We therefore call \mathcal{A}_N the *Artin fan* of such small X . By the construction, the Artin fan of X is functorial with respect to strict morphisms.

Now consider a logarithmic algebraic stack X with a groupoid presentation

$$V \rightrightarrows U \rightarrow X$$

in which U and V are disjoint unions of *small* Zariski logarithmic schemes. Then V and U have Artin fans \mathcal{V} and \mathcal{U} and we obtain strict morphisms of Artin fans $\mathcal{V} \rightarrow \mathcal{U}$. Strict morphisms of Artin fans are étale, so this is a diagram of étale spaces over **Log**. It therefore has a colimit, also an étale space over **Log**, which we call the Artin fan of X .

5.4 Functoriality of Artin Fans: Problem and Fix

The universal property of the Artin fan implies immediately that Artin fans are functorial with respect to *strict* morphisms of logarithmic schemes. They are not functorial in general, but we will be able to salvage a weak replacement for functoriality in which morphisms of logarithmic schemes induce correspondences of Artin fans.

5.4.1 The Failure of Functoriality

We use the notation for the punctured Whitney umbrella introduced in Sect. 4.6.2. As X has a global chart, its Artin fan \mathcal{X} is easily seen to be \mathcal{A}^2 . The Artin fan \mathcal{Y} of Y is the quotient of \mathcal{A}^2 by the action of $\mathbb{Z}/2\mathbb{Z}$ exchanging the components, *as a representable étale space over Log*. In other words, the group action induces an étale equivalence relative to **Log** by taking the image of the action map

$$\mathbb{Z}/2\mathbb{Z} \times \mathcal{A}^2 \rightarrow \mathcal{A}^2 \times_{\text{Log}} \mathcal{A}^2.$$

A logarithmic morphism from a logarithmic scheme S into $\mathcal{A}^2 \times_{\text{Log}} \mathcal{A}^2$ consists of two maps $\mathbb{N}^2 \rightarrow \Gamma(S, \overline{M}_S)$ and an isomorphism between the induced logarithmic structures M_1 and M_2 commuting with the projection to M_S . This implies that

$$\mathcal{A}^2 \times_{\text{Log}} \mathcal{A}^2 = \mathcal{A}^2 \amalg_{\mathcal{A}^0} \mathcal{A}^2$$

where by \mathcal{A}^0 we mean the open point of \mathcal{A}^2 . The nontrivial projection in

$$\mathcal{A}^2 \amalg_{\mathcal{A}^0} \mathcal{A}^2 = \mathcal{A}^2 \times_{\text{Log}} \mathcal{A}^2 \rightrightarrows \mathcal{A}^2$$

is given by the identity map on one component and the exchange of coordinates on the other component. (The trivial projection is the identity on both components.)

It is now easy to see that $\mathbb{Z}/2\mathbb{Z} \times \mathcal{A}^2$ surjects onto $\mathcal{A}^2 \amalg_{\mathcal{A}^0} \mathcal{A}^2$, so the Artin fan of Y is the image of \mathcal{A}^2 in **Log**. We will write $\mathcal{A}^{[2]}$ for this open substack and observe that it represents the functor sending a logarithmic scheme S to the category of pairs (\overline{M}, φ) where \overline{M} is an étale sheaf of monoids on S and $\varphi : \overline{M} \rightarrow \overline{M}_S$ is a strict morphism *that can be presented étale locally by a map* $\mathbb{N}^2 \rightarrow \overline{M}_S$. Equivalently, it is a logarithmic structure over M_S that has étale-local charts by \mathbb{N}^2 .

Of course, there is a map $\mathcal{X} \rightarrow \mathcal{Y}$ consistent with the maps $X \rightarrow Y$: it is the canonical projection $\mathcal{A}^2 \rightarrow \mathcal{A}^{[2]}$. However, we take \widetilde{X} to be the logarithmic blowup of X at $\{0\} \times \mathbb{G}_m$ and let \widetilde{Y} be the logarithmic blowup of Y at the image of this locus. Then we have a Cartesian diagram, since X is flat over Y :

$$\begin{array}{ccc} \tilde{X} & \longrightarrow & \tilde{Y} \\ \downarrow & & \downarrow \\ X & \longrightarrow & Y \end{array}$$

We will compute the Artin fans of \tilde{X} and \tilde{Y} .

Since X and Y are flat over their Artin fans, the blowups \tilde{X} and \tilde{Y} are pulled back from the blowups $\tilde{\mathcal{X}}$ and $\tilde{\mathcal{Y}}$ of \mathcal{X} and \mathcal{Y} . Furthermore, the square in the diagram below is Cartesian:

$$\begin{array}{ccccc} \tilde{\mathcal{X}} & \longrightarrow & \tilde{\mathcal{Y}} & \longrightarrow & \mathcal{Z} \\ \downarrow & & \downarrow & \swarrow \text{dashed} & \\ \mathcal{X} & \longrightarrow & \mathcal{Y} & & \end{array}$$

We have written \mathcal{Z} for the Artin fan of \tilde{Y} . Since $\tilde{Y} \rightarrow \tilde{\mathcal{Y}}$ is strict and smooth with connected fibers, the map $\tilde{Y} \rightarrow \mathcal{Z}$ factors through $\tilde{\mathcal{Y}}$. We will see in a moment that there is no dashed arrow making the triangle on the right above commutative, thus witnessing the failure of the functoriality of the Artin fan construction with respect to the morphism $\tilde{Y} \rightarrow Y$.³

The reason no map $\mathcal{Z} \rightarrow \mathcal{Y}$ can exist making the diagram above commutative is that $\tilde{\mathcal{Y}}$ has monodromy at the generic point of its exceptional divisor, pulled back from the monodromy at the closed point of \mathcal{Y} . However, the image of the exceptional divisor of $\tilde{\mathcal{Y}}$ in \mathcal{Z} is a divisor, and $\mathcal{Z} \rightarrow \mathbf{Log}$ is representable by algebraic spaces. No rank 1 logarithmic structure can have monodromy, so there is no monodromy at the image of the exceptional divisor in \mathcal{Z} .

A variant of the last paragraph shows that there is no commutative diagram

$$\begin{array}{ccc} \tilde{Y} & \longrightarrow & \mathcal{Z} \\ \downarrow & & \downarrow \\ Y & \longrightarrow & \mathcal{Y} \end{array}$$

We can find a loop γ in the exceptional divisor E of \tilde{Y} that projects to a nontrivial loop in Y , around which the logarithmic structure of Y has nontrivial monodromy. Even though the logarithmic structure of \tilde{Y} has no monodromy around γ , the image

³More specifically, there is no way to make the Artin fan functorial while also making the maps $Y \rightarrow \mathcal{Y}$ natural in Y .

of γ in \mathcal{Y} is nontrivial. But all of the exceptional divisor E is collapsed to a point in \mathcal{Z} , so the image of γ in \mathcal{Y} must act trivially.⁴

5.4.2 The Patch

There seem to be two ways to get around this failure of functoriality. The first is to allow the Artin fan to include more information about the fundamental group of the original logarithmic scheme X . However, the most naive application of this principle would introduce the entire étale homotopy type of X into the Artin fan, sacrificing Artin fans' essentially combinatorial nature.

Another approach draws inspiration from Olsson's stacks of diagrams of logarithmic structures [47]. Let $\mathbf{Log}^{[1]}$ be the stack whose S -points are morphisms of logarithmic structures $M_1 \rightarrow M_2$ on S . A morphism of logarithmic schemes $X \rightarrow Y$ induces a commutative diagram

$$\begin{array}{ccc} X & \longrightarrow & \mathbf{Log}^{[1]} \\ \downarrow & & \downarrow \\ Y & \longrightarrow & \mathbf{Log} \end{array}$$

where $\mathbf{Log}^{[1]} \rightarrow \mathbf{Log}$ sends $M_1 \rightarrow M_2$ to M_1 . If $\mathbf{Log}^{[1]}$ is given M_2 as its logarithmic structure, this is a commutative diagram of logarithmic algebraic stacks. The construction of the Artin fan works in a relative situation, and we take $\mathcal{Y} = \pi_0(Y/\mathbf{Log})$ and $\mathcal{X} = \pi_0(X/\mathbf{Log}^{[1]})$. Note that $\mathcal{Y} \times_{\mathbf{Log}} \mathbf{Log}^{[1]}$ is étale over $\mathbf{Log}^{[1]}$, so we get a map

$$\mathcal{X} \rightarrow \mathcal{Y} \times_{\mathbf{Log}} \mathbf{Log}^{[1]} \rightarrow \mathcal{Y}$$

salvaging a commutative diagram:

Theorem 5.7 ([3, Corollary 3.3.5]). *For any morphism of logarithmic schemes $X \rightarrow Y$ with étale-locally connected logarithmic strata there is an initial commutative diagram*

⁴Here is a rigorous version of the above argument. Observe that \mathcal{Y} is the quotient of $\mathcal{R} \rightrightarrows \mathcal{X}$, where $\mathcal{X} \simeq \mathcal{A}^2$ and \mathcal{R} is two copies of \mathcal{X} joined along their open point. Pulling \mathcal{X} back to the center of the blowup $Y_0 \simeq \mathbb{G}_m$ inside Y yields the cover by $\mathbb{G}_m \times \{0\} \subset X$. This cover has nontrivial monodromy. On the other hand, $\mathcal{Z} \simeq \mathcal{A}^2$ has no nontrivial étale covers. Therefore the pullback of $\mathcal{X} \times_{\mathcal{Y}} \mathcal{Z}$ via $\tilde{Y} \rightarrow \mathcal{Z}$ yields a trivial étale cover of Y_0 .

$$\begin{array}{ccccc}
 X & \longrightarrow & \mathcal{X} & \longrightarrow & \mathbf{Log}^{[1]} \\
 \downarrow & & \downarrow & & \downarrow \\
 Y & \longrightarrow & \mathcal{Y} & \longrightarrow & \mathbf{Log}
 \end{array}$$

in which the horizontal arrows are strict and both \mathcal{X} and \mathcal{Y} are Artin fans representable by algebraic spaces relative to $\mathbf{Log}^{[1]}$ and \mathbf{Log} , respectively.

6 Algebraic Applications of Artin Fans

6.1 Gromov–Witten Theory and Relative Gromov–Witten Theory

Algebraic Gromov–Witten theory is the study of the virtually enumerative invariants, known as Gromov–Witten invariants, of algebraic curves on a smooth target variety X . In Gromov–Witten theory one integrates cohomology classes on X against the virtual fundamental class $[\overline{\mathcal{M}}_\Gamma(X)]^{\text{vir}}$ of the moduli space $\overline{\mathcal{M}}_\Gamma(X)$ of stable maps with target X . The subscript Γ indicates fixed numerical invariants, including the genus of the domain curve, the number of marked points on it, and the homology class of its image.

Relative Gromov–Witten theory comes from efforts to define Gromov–Witten invariants for degenerations of complicated targets that, while singular, are still geometrically simple. In the mild setting of two smooth varieties meeting along a smooth divisor, such a theory has been developed by Li [37, 38], following work in symplectic geometry by Li–Ruan [39] and Ionel–Parker [30, 31].

Working over \mathbb{C} , one considers a degeneration

$$\begin{array}{ccc}
 Y_1 \sqcup_D Y_2 & \subset & X \\
 \downarrow & & \downarrow \pi \\
 b_0 & \subset & B
 \end{array} \tag{2}$$

where $\pi : X \rightarrow B$ is a flat, projective morphism from a smooth variety to a smooth curve, and where the singular fiber $X_0 = Y_1 \sqcup_D Y_2$ consists of two smooth varieties meeting along a smooth divisor.

Jun Li proved an algebro-geometric degeneration formula through which one can recover Gromov–Witten invariants of the possibly complicated but smooth general fiber of X from *relative Gromov–Witten invariants* determined by a space $\overline{\mathcal{M}}_\Gamma(Y_i)$ of *relative stable maps* to each of the two smooth components Y_i of X_0 . Here, in

addition to the genus, the number of markings, and homology class, one must fix the *contact orders* of the curve with the given divisor.

6.1.1 Expanded Degenerations and Pairs

In order to define Gromov–Witten invariants of the *singular* degenerate fiber X_0 , Jun Li constructed a whole family of *expansions* $X'_0 \rightarrow X_0$, where

$$X'_0 = Y_1 \sqcup_D P \sqcup_D \cdots \sqcup_D P \sqcup_D Y_2.$$

Here P is the projective completion of $N_{D/Y_1} \simeq N_{D/Y_2}^\vee$ (explicitly, it is $\mathbb{P}(\mathcal{O} \oplus N_{D/Y_1})$) and the gluing over D attaches 0-sections to ∞ -sections.

Similarly, in order to guarantee that contact orders of maps in each Y_i are maintained, Jun Li constructed a family of expansions $Y'_i \rightarrow Y_i$ where

$$Y'_i = Y_i \sqcup_D P \sqcup_D \cdots \sqcup_D P.$$

Here is the first point where Artin fans, in their simplest form and even without logarithmic structures, become of use:

In Jun Li’s construction, not every deformation of an expansion $Y'_i \rightarrow Y_i$ is itself an expansion. For instance, the expansion $Y_i \sqcup_D P$ can deform to $Y_i \sqcup_D P'$, where $P' = \mathbb{P}(\mathcal{O} \oplus N'')$ with N'' a deformation of N_{D/Y_1} . Precisely the same problem occurs with deformation of an expansion $X'_0 \rightarrow X_0$.

6.1.2 The Artin Fan as the Universal Target, and Its Expansions

In [5], following ideas in [14], it was noted that the Y_i has a canonical map $Y_i \rightarrow \mathcal{A} := [\mathbb{A}^1/\mathbb{G}_m]$. From the point of view of the present text, $\mathcal{A} = \mathcal{A}_{Y_i}$, the Artin fan associated to the divisorial logarithmic structure (Y_i, D) , and its divisor $\mathcal{D} = B\mathbb{G}_m \subset \mathcal{A}$ is the universal divisor. Next, if $\mathcal{A}' \rightarrow \mathcal{A}$ is an expansion, then *all deformations of $\mathcal{A}' \rightarrow \mathcal{A}$ are expansions of \mathcal{A} in the sense of Jun Li*. And finally, any expansion $Y'_i \rightarrow Y_i$ is obtained as the pullback $Y'_i = \mathcal{A}' \times_{\mathcal{A}} Y_i$ of some expansion $\mathcal{A}' \rightarrow \mathcal{A}$.

This means that the moduli space of expansions of any pair (Y_i, D) is identical to the moduli space of expansions of $(\mathcal{A}, \mathcal{D})$, and the expansions themselves are obtained by pullback.

A similar picture occurs for degenerations: the Artin fan of $X \rightarrow B$ is the morphism $\mathcal{A} \times \mathcal{A} \rightarrow \mathcal{A}$ induced by the multiplication morphism $\mathbb{A}^2 \rightarrow \mathbb{A}^1$:

$$\begin{array}{ccc} X & \longrightarrow & \mathcal{A}^2 \\ \downarrow & & \downarrow \\ B & \longrightarrow & \mathcal{A}. \end{array}$$

There is again a stack of universal expansions of $\mathcal{A} \times \mathcal{A} \rightarrow \mathcal{A}$, and every expansion $X'_0 \rightarrow X_0$ is the pullback of a fiber of the universal expansion:

$$\begin{array}{ccc} X'_0 & \longrightarrow & (\mathcal{A}^2)' \\ \downarrow & & \downarrow \\ X_0 & \longrightarrow & \mathcal{A}^2 \times_A B \end{array}$$

From the point of view of logarithmic geometry this is not surprising: expansions are always stable under *logarithmic* deformations. But the approach through Artin fans provides us with further results, which we outline below.

6.1.3 Redefining Obstructions

Using expanded degenerations and expanded pairs, Jun Li defined moduli spaces of *degenerate stable maps* $\overline{\mathcal{M}}_\Gamma(X/B)$ and of *relative stable maps* $\overline{\mathcal{M}}_\Gamma(Y_i, D)$. Jun Li had an additional challenge in defining the virtual fundamental classes of these spaces, which he constructed by bare hands. With a little bit of hindsight, we now know that Li’s virtual fundamental classes are associated to the natural *relative obstruction theories* of the morphisms $\overline{\mathcal{M}}_\Gamma(X/B) \rightarrow \mathfrak{M}_\Gamma(\mathcal{A}^2/\mathcal{A})$ and $\overline{\mathcal{M}}_\Gamma(Y_i, D) \rightarrow \mathfrak{M}_\Gamma(\mathcal{A}, \mathcal{D})$, where $\mathfrak{M}_\Gamma(\mathcal{A}^2/\mathcal{A})$ and $\mathfrak{M}_\Gamma(\mathcal{A}, \mathcal{D})$ are the associated moduli spaces of *prestable maps*. Moreover, the virtual fundamental classes can be understood with machinery available off the shelf of any deformation theory emporium. This observation from [7] made it possible to prove a number of comparison results, including those described below.

6.1.4 Other Approaches and Comparison Theorems

Denote by (Y, D) a smooth pair, consisting of a smooth projective variety Y with a smooth and irreducible divisor D . Jun Li’s moduli space $\overline{\mathcal{M}}_\Gamma(Y, D)$ of relative stable maps to (Y, D) provides an algebraic setting for relative Gromov–Witten theory. Only recently have efforts to generalize the theory to more complicated singular targets come to fruition [1, 24, 29, 48, 49]. Because of the technical difficulty of Jun Li’s approach, several alternate approaches to the relative Gromov–Witten invariants of (Y, D) have been developed:

- $\text{Li}_\Gamma^{\text{stab}}(Y, D)$: Li’s original moduli space of relative stable maps;
- $\text{AF}_\Gamma^{\text{stab}}(Y, D)$: Abramovich–Fantechi’s stable orbifold maps with expansions [2];
- $\text{Kim}_\Gamma^{\text{stab}}(Y, D)$: Kim’s stable logarithmic maps with expansions [36]; and
- $\text{ACGS}_\Gamma^{\text{stab}}(Y, D)$: Abramovich–Chen and Gross–Siebert’s stable logarithmic maps without expansions [1, 17, 24].


There are analogous constructions

$$\mathrm{Li}_\Gamma^{\mathrm{stab}}(X/B) \quad \mathrm{AF}_\Gamma^{\mathrm{stab}}(X/B) \quad \mathrm{Kim}_\Gamma^{\mathrm{stab}}(X/B) \quad \mathrm{ACGS}_\Gamma^{\mathrm{stab}}(X/B)$$

for a degeneration.

This poses a new conundrum: how do these approaches compare? The answer, which depends on the Artin fan \mathcal{A} , is as follows:

Theorem 6.1 ([7, Theorem 1.1]). *There are maps*



such that

$$\Psi_*[\mathrm{AF}]^{\mathrm{vir}} = [\mathrm{Li}]^{\mathrm{vir}} \quad \Theta_*[\mathrm{Kim}]^{\mathrm{vir}} = [\mathrm{Li}]^{\mathrm{vir}} \quad \Upsilon_*[\mathrm{Kim}]^{\mathrm{vir}} = [\mathrm{ACGS}]^{\mathrm{vir}}.$$

In particular, the Gromov–Witten invariants associated to these four theories coincide.

The principle behind this comparison for each of the three maps Ψ , Θ , and Υ is the same, and we illustrate it on Ψ : there are algebraic stacks parametrizing orbifold and relative stable maps to the Artin fan $(\mathcal{A}, \mathcal{D})$, which we denote here by $\mathrm{AF}_\Gamma(\mathcal{A}, \mathcal{D})$ and $\mathrm{Li}_\Gamma(\mathcal{A}, \mathcal{D})$. These stacks sit in a Cartesian diagram

$$\begin{array}{ccc}
 \mathrm{AF}_\Gamma^{\mathrm{stab}}(Y, D) & \xrightarrow{\Psi} & \mathrm{Li}_\Gamma^{\mathrm{stab}}(Y, D) \\
 \pi_{\mathrm{AF}} \downarrow & & \downarrow \pi_{\mathrm{Li}} \\
 \mathrm{AF}_\Gamma(\mathcal{A}, \mathcal{D}) & \xrightarrow{\Psi_{\mathcal{A}}} & \mathrm{Li}_\Gamma(\mathcal{A}, \mathcal{D})
 \end{array}$$

such that

- (1) The virtual fundamental classes of the spaces $\mathrm{AF}_\Gamma^{\mathrm{stab}}(Y, D)$ and $\mathrm{Li}_\Gamma^{\mathrm{stab}}(Y, D)$ may be computed using the natural obstruction theory relative to the vertical arrows π_{AF} and π_{Li} .
- (2) The obstruction theory for π_{AF} is the pullback of the obstruction theory of π_{Li} .
- (3) The morphism $\Psi_{\mathcal{A}}$ is birational.

One then applies a general comparison result of Costello [18, Theorem 5.0.1] to obtain the theorem.

What makes all this possible is the fact that the virtual fundamental classes of $\mathrm{AF}_\Gamma(\mathcal{A}, \mathcal{D})$ and $\mathrm{Li}_\Gamma(\mathcal{A}, \mathcal{D})$ agree with their fundamental classes. This in turn results from the fact that the Artin fan $(\mathcal{A}, \mathcal{D})$ of (X, D) discards all the complicated

geometry of (X, D) , retaining just enough algebraic structure to afford stacks of maps such as $\mathbf{AF}_\Gamma(\mathcal{A}, \mathcal{D})$ and $\mathbf{Li}_\Gamma(\mathcal{A}, \mathcal{D})$. In effect, the *virtual* birationality of Ψ is due to the *genuine* birationality of $\Psi_{\mathcal{A}}$.

Similar results were obtained by similar methods in [16] and [4].

6.2 Birational Invariance for Logarithmic Stable Maps

More general Artin fans have found applications in the logarithmic approach to Gromov–Witten theory.

Let X be a projective *logarithmically smooth* variety over \mathbb{C} , and denote by $\overline{\mathcal{M}}(X)$ the moduli space of *logarithmic stable maps* to X , as defined in [1, 17, 24]. Fix a logarithmically étale morphism $h : Y \rightarrow X$ of projective logarithmically smooth varieties. There is a natural morphism $\overline{\mathcal{M}}(h) : \overline{\mathcal{M}}(Y) \rightarrow \overline{\mathcal{M}}(X)$ induced by h , and the central result of [3] is the following pushforward statement for virtual classes.

Theorem 6.2. $\overline{\mathcal{M}}(h)_*([\overline{\mathcal{M}}(Y)]^{\text{vir}}) = [\overline{\mathcal{M}}(X)]^{\text{vir}}$. In particular the logarithmic Gromov–Witten invariants of X and Y coincide.

This theorem is proven by working relative to the underlying morphism of Artin fans $\mathcal{Y} \rightarrow \mathcal{X}$ provided by Theorem 5.7. This morphism and h fit into a Cartesian diagram

$$\begin{array}{ccc}
 Y & \xrightarrow{h} & X \\
 \downarrow & & \downarrow \\
 \mathcal{Y} & \longrightarrow & \mathcal{X}.
 \end{array} \tag{3}$$

One carefully constructs a Cartesian diagram of moduli spaces

$$\begin{array}{ccc}
 \overline{\mathcal{M}}(Y) & \xrightarrow{\overline{\mathcal{M}}(h)} & \overline{\mathcal{M}}(X) \\
 \downarrow & & \downarrow \\
 \overline{\mathcal{M}}(\mathcal{Y}) & \longrightarrow & \overline{\mathcal{M}}(\mathcal{X})
 \end{array} \tag{4}$$

to which principles (1), (2), and (3) in the proof of Theorem 6.1 apply. Costello’s comparison theorem [18, Theorem 5.0.1] again gives the result.

6.3 Boundedness of Logarithmic Stable Maps

A recent application of both Theorem 6.2 and the theory of Artin fans can be found in [8], where a general statement for the boundedness of logarithmic stable maps to projective logarithmic schemes is proven.

Theorem 6.3 ([8, Theorem 1.1.1]). *Let X be a projective logarithmic scheme. Then the stack $\overline{\mathcal{M}}_\Gamma(X)$ of stable logarithmic maps to X with discrete data Γ is of finite type.*

The boundedness of $\overline{\mathcal{M}}_\Gamma(X)$ has been established when the characteristic monoid \overline{M}_X is globally generated in [1, 17], and more generally when the associated group $\overline{M}_X^{\text{gp}}$ is globally generated in [24]. The strategy used in [8] for the general setting is to reduce to the case of a globally generated sheaf of monoids \overline{M}_X by studying the behavior of stable logarithmic maps under an appropriate modification of X . This is accomplished by modifying the Artin fan $\mathcal{X} = \mathcal{A}_X$ constructed in Sect. 5.3, lifting this modification to the level of logarithmic schemes, and applying a virtual birationality result refining Theorem 6.2.

A key step is the modification of \mathcal{X} :

Proposition 6.4 ([8, Proposition 1.3.1]). *Let \mathcal{X} be an Artin fan. Then there exists a projective, birational, and logarithmically étale morphism $\mathcal{Y} \rightarrow \mathcal{X}$ such that \mathcal{Y} is a smooth Artin fan, and the characteristic sheaf $\overline{M}_\mathcal{Y}$ is globally generated and locally free.*

The proof of this Proposition is analogous to [35, I.11]: successive star subdivisions are applied until each cone is smooth (see Example 4.10), and barycentric subdivisions guarantee that the resulting logarithmic structure has no monodromy.

7 Skeletons and Tropicalization

7.1 Berkovich Spaces

Ever since [56] it has been known that affinoid algebras are the correct coordinate rings for defining non-Archimedean analogues of complex analytic spaces. Working just with affinoid algebras as coordinate rings, we have enough information to build an intricate theory with many applications. The work of Berkovich [11] and [12] beautifully enriches this theory by providing us with an alternative definition of non-Archimedean analytic spaces, which naturally come with underlying topological spaces that have many of the favorable properties of complex analytic spaces, such as being locally path-connected and locally compact.

Let \mathbf{k} be a non-Archimedean field, i.e., suppose that \mathbf{k} is complete with respect to a non-Archimedean absolute value $|\cdot|$. We explicitly allow \mathbf{k} to carry the trivial absolute value. If $U = \text{Spec } A$ is an affine scheme of finite type over \mathbf{k} , as a set the analytic space U^{an} associated to U is equal to the set of multiplicative seminorms on

A that restrict to the given absolute value on \mathbf{k} . We usually write x for a point in U^{an} and $|\cdot|_x$, if we want to emphasize that x is thought of as a multiplicative seminorm on A . The topology on U^{an} is the coarsest that makes the maps

$$\begin{aligned}
 U^{an} &\longrightarrow \mathbb{R} \\
 x &\longmapsto |f|_x
 \end{aligned}$$

continuous for all $f \in A$. There is a natural continuous *structure morphism* $\rho : U^{an} \rightarrow U$ given by sending $x \in U^{an}$ to the preimage of zero $\{f \in A \mid |f|_x = 0\}$. A morphism $f : U \rightarrow V$ between affine schemes $U = \text{Spec } A$ and $V = \text{Spec } B$ of finite type over \mathbf{k} , given by a \mathbf{k} -algebra homomorphism $f^\# : B \rightarrow A$, induces a natural continuous map $f^{an} : U^{an} \rightarrow V^{an}$ given by associating to $x \in U^{an}$ the point $f(x) \in V^{an}$ given as the multiplicative seminorm

$$|\cdot|_{f(x)} = |\cdot|_x \circ f^\#$$

on B . The association $f \mapsto f^{an}$ is functorial in f .

Let X be a scheme that is locally of finite type over \mathbf{k} . Choose a covering $U_i = \text{Spec } A_i$ of X by open affines. Then the *analytic space* X^{an} associated to X is defined by glueing the U_i^{an} over the open subsets $\rho^{-1}(U_i \cap U_j)$. It is easy to see that this construction does not depend on the choice of a covering and that it is functorial with respect to morphisms of schemes over \mathbf{k} . We refrain from describing the structure sheaf on X^{an} , since we are only going to be interested in the topological properties of X^{an} , and refer the reader to [11, 12, 57] for further details.

Example 7.1. Suppose that \mathbf{k} is algebraically closed and endowed with the trivial absolute value. Let t be a coordinate on the affine line $\mathbb{A}^1 = \text{Spec } \mathbf{k}[t]$. One can classify the points in $(\mathbb{A}^1)^{an}$ as follows:

- For every $a \in \mathbf{k}$ and $r \in [0, 1)$ we have the seminorm $|\cdot|_{a,r}$ on $\mathbf{k}[t]$ that is uniquely determined by

$$|t - a|_{a,r} = r,$$

and

- for $r \in [1, \infty)$ we have the seminorm $|\cdot|_{\infty,r}$ that is uniquely determined by


$$|t - a|_{\infty,r} = r$$

for all $a \in \mathbf{k}$.

Noting that

$$\lim_{r \rightarrow 1} |\cdot|_{a,r} = |\cdot|_{\infty,1}$$

for all $a \in \mathbf{k}$ we can visualize $(\mathbb{A}^1)^{an}$ as follows:



In particular, we can embed the closed points of $\mathbb{A}^1(\mathbf{k})$ into $(\mathbb{A}^1)^{an}$ by the association $a \mapsto |\cdot|_{a,0}$ for $a \in \mathbf{k}$.

One can give an alternative description of X^{an} as the set of equivalence classes of non-Archimedean points of X . A *non-Archimedean point* of X consists of a pair (K, ϕ) where K is a non-Archimedean extension of \mathbf{k} and $\phi : \text{Spec } K \rightarrow X$. Two non-Archimedean points (K, ϕ) and (L, ψ) of X are *equivalent*, if there is a common non-Archimedean extension Ω of both K and L such that the diagram commutes.

$$\begin{array}{ccc}
 \text{Spec } \Omega & \longrightarrow & \text{Spec } L \\
 \downarrow & & \downarrow \psi \\
 \text{Spec } K & \xrightarrow{\phi} & X
 \end{array}$$

Proposition 7.2. *The analytic space X^{an} is equal to the set of non-Archimedean points modulo equivalence.*

Proof. We may assume that $X = \text{Spec } A$ is affine. A non-Archimedean point (K, ϕ) of X naturally induces a multiplicative seminorm

$$A \xrightarrow{\phi^\#} K \xrightarrow{|\cdot|} \mathbb{R}$$

on A . Conversely, given $x \in X^{an}$, we can consider the integral domain $A / \ker |\cdot|_x$ and form the completion $\mathcal{H}(x)$ of its field of fractions. The natural homomorphism $A \rightarrow \mathcal{H}(x)$ defines a non-Archimedean point on X and these two constructions are inverse up to equivalence. □

Now suppose that \mathbf{k} is endowed with the trivial absolute value. In [58, Sect. 1] Thuillier introduces a slight variant of the analytification functor that functorially


associates to a scheme locally of finite type over \mathbf{k} a non-Archimedean analytic space X^\natural . Its points are pairs (R, ϕ) consisting of a valuation ring R extending \mathbf{k} together with a morphism $\phi : \text{Spec } R \rightarrow X$ modulo an equivalence relation as above: two pairs (R, ϕ) and (S, ψ) are *equivalent*, if there is a common valuation ring \mathcal{O} extending both R and S such that the diagram commutes.

$$\begin{array}{ccc} \text{Spec } \mathcal{O} & \longrightarrow & \text{Spec } S \\ \downarrow & & \downarrow \psi \\ \text{Spec } R & \xrightarrow{\phi} & X \end{array}$$

So, if $U = \text{Spec } A$ is affine, then U^\natural is nothing but the set of multiplicative seminorms $|\cdot|_x$ on A that are bounded, i.e., that fulfill $|f|_x \leq 1$ for all $f \in A$. The topology on U^\natural is the one induced from U^{an} and there is a natural anti-continuous *reduction map* $r : U^\natural \rightarrow U$ that sends $x \in U^\natural$ to the prime ideal $\{f \in A \mid |f|_x < 1\}$. For a general scheme X locally of finite type over \mathbf{k} , we again choose a covering U_i by open affine subsets, and now glue the U_i^\natural over the closed subsets $r^{-1}(U_i \cap U_j)$ in order to obtain X^\natural .

As an immediate application of the valuative criteria for separatedness and properness, we obtain that, if X is separated, then X^\natural is naturally a locally closed subspace of X^{an} , and, if X is complete, then $X^\natural = X^{an}$.

Example 7.3. The space $(\mathbb{A}^1)^\natural$ is precisely the subspace of $(\mathbb{A}^1)^{an}$ consisting of the points $|\cdot|_{a,r}$ for $a \in \mathbf{k}$ and $r \in [0, 1)$ as well as the Gauss point $|\cdot|_{\infty,1}$.



7.2 The Case of Toric Varieties

Let \mathbf{k} be a non-Archimedean field. As observed in [10, 22, 25, 26], there is an intricate relationship between non-Archimedean analytic geometry and tropical geometry. In particular, in many interesting situations the tropicalization of an

algebraic variety X over \mathbf{k} can be regarded as a natural deformation retract of X^{an} , a so-called *skeleton* of X^{an} . In this section we are going to give a detailed explanation of this relationship in the simplest possible case, that of toric varieties.

Let T be a split algebraic torus with character lattice M and co-character lattice N , and $X = X(\Delta)$ a T -toric variety that is defined by a rational polyhedral fan Δ in $N_{\mathbb{R}}$. We refer the reader to Sect. 2 for a brief summary of this beautiful theory, and to [23] and [19] for a more thorough account.

In [32] and [50] Kajiwara and Payne independently construct a *tropicalization map*

$$\text{trop}_{\Delta} : X^{an} \longrightarrow N_{\mathbb{R}}(\Delta)$$

associated to X , whose codomain is a partial compactification of $N_{\mathbb{R}}$, uniquely determined by Δ (also see [51, Sect. 4] and [53, Sect. 3]).

For a cone σ in Δ set $N_{\mathbb{R}}(\sigma) = \text{Hom}(S_{\sigma}, \overline{\mathbb{R}})$, where S_{σ} denotes the toric monoid $\sigma^{\vee} \cap M$ and write $\overline{\mathbb{R}}$ for the additive monoid $(\mathbb{R} \sqcup \{\infty\}, +)$. Endow $N_{\mathbb{R}}(\sigma)$ with the topology of pointwise convergence.

Lemma 7.4 ([53, Proposition 3.4]).

- (i) *The space $N_{\mathbb{R}}(\sigma)$ has a stratification by locally closed subsets isomorphic to the vector spaces $N_{\mathbb{R}}/\text{Span}(\tau)$ for all faces τ of σ .*
- (ii) *For a face τ of σ the natural map $S_{\sigma} \rightarrow S_{\tau}$ induces the open embedding*

$$N_{\mathbb{R}}(\tau) \hookrightarrow N_{\mathbb{R}}(\sigma)$$

that identifies $N_{\mathbb{R}}(\tau)$ with the union of strata in $N_{\mathbb{R}}(\sigma)$ corresponding to faces of τ in $N_{\mathbb{R}}(\sigma)$.

So one can think of $N_{\mathbb{R}}(\sigma)$ as a partial compactification of $N_{\mathbb{R}}$ given by adding a vector space $N_{\mathbb{R}}/\text{Span}(\tau)$ at infinity for every face $\tau \neq 0$ of σ . The partial compactification $N_{\mathbb{R}}(\Delta)$ of $N_{\mathbb{R}}$ is defined to be the colimit of the $N_{\mathbb{R}}(\sigma)$ for all cones σ in Δ . Since the stratifications on the $N_{\mathbb{R}}(\sigma)$ are compatible, the space $N_{\mathbb{R}}(\Delta)$ is a partial compactification of $N_{\mathbb{R}}$ that carries a stratification by locally closed subsets isomorphic to $N_{\mathbb{R}}/\text{Span}(\sigma)$ for every cone σ in Δ .

On the T -invariant open affine subset $X_{\sigma} = \text{Spec } \mathbf{k}[S_{\sigma}]$ the *tropicalization map*

$$\text{trop}_{\sigma} : X_{\sigma}^{an} \longrightarrow N_{\mathbb{R}}(\sigma)$$

is defined by associating to an element $x \in X^{an}$ the homomorphism $s \mapsto -\log |\chi^s|_x$ in $N_{\mathbb{R}}(\sigma) = \text{Hom}(S_{\sigma}, \overline{\mathbb{R}})$.

Lemma 7.5. *The tropicalization map trop_{σ} is continuous and, for a face τ of σ , the natural diagram*

$$\begin{array}{ccc}
 X_\tau^{an} & \xrightarrow{\text{trop}_\tau} & N_{\mathbb{R}}(\tau) \\
 \subseteq \downarrow & & \downarrow \subseteq \\
 X_\sigma^{an} & \xrightarrow{\text{trop}_\sigma} & N_{\mathbb{R}}(\sigma)
 \end{array}$$

commutes and is Cartesian.

Therefore we can glue the trop_σ on local T -invariant patches X_σ and obtain a global continuous tropicalization map

$$\text{trop}_\Delta : X^{an} \longrightarrow N_{\mathbb{R}}(\Delta)$$

that restricts to trop_σ on T -invariant open affine subsets X_σ . Its restriction to T -orbits is the usual tropicalization map in the sense of [26, Sect. 3].

The following Proposition 7.6 is well-known among experts and can be found in [58, Sect. 2] in the constant coefficient case, i.e., the case that \mathbf{k} is trivially valued.

Proposition 7.6. *The tropicalization map $\text{trop}_\Delta : X^{an} \rightarrow N_{\mathbb{R}}(\Delta)$ has a continuous section J_Δ and the composition*

$$\mathbf{p}_\Delta = J_\Delta \circ \text{trop}_\Delta : X^{an} \rightarrow X^{an}$$

defines a strong deformation retraction.

The deformation retract

$$\mathfrak{S}(X) = J_\Delta(N_{\mathbb{R}}(\Delta)) = \mathbf{p}_\Delta(X^{an})$$

is said to be the *non-Archimedean skeleton* of X^{an} .

Proof Sketch of Proposition 7.6. Consider a T -invariant open affine subset $X_\sigma = \text{Spec } \mathbf{k}[S_\sigma]$. We may construct the section $J_\sigma : N_{\mathbb{R}}(\sigma) \rightarrow X_\sigma^{an}$ by associating to $u \in N_{\mathbb{R}}(\sigma) = \text{Hom}(S_\sigma, \overline{\mathbb{R}})$ the seminorm $J_\sigma(u)$ defined by

$$J_\sigma(u)(f) = \max_{s \in S_\sigma} \left\{ |a_s| \exp(-u(s)) \right\}$$

for $f = \sum_s a_s \chi^s \in \mathbf{k}[S_\sigma]$. A direct verification shows that J_σ is continuous and fulfills $\text{trop}_\sigma \circ J_\sigma = \text{id}_{N_{\mathbb{R}}(\sigma)}$.

The construction of J_σ is compatible with restrictions to T -invariant affine open subsets and we obtain a global section $J_\Delta : N_{\mathbb{R}}(\Delta) \rightarrow X^{an}$ of the tropicalization map $\text{trop}_\Delta : X^{an} \rightarrow N_{\mathbb{R}}(\Delta)$.

Since J_Δ is a section of trop_Δ , the continuous map

$$\mathbf{p}_\Delta = J_\Delta \circ \text{trop}_\Delta : X^{an} \longrightarrow X^{an}$$

is a retraction map. On X_σ the image $\mathbf{p}_\sigma(x)$ of $x \in X_\sigma^{an}$ is the seminorm given by

$$\mathbf{p}_\sigma(x)(f) = \max_{s \in S_\sigma} \{ |a_s| |\chi^s|_x \}$$

for $f = \sum_s a_s \chi^s \in \mathbf{k}[S_\sigma]$. The arguments in [58, Sect. 2.2] generalize to this situation and show that \mathbf{p}_Δ is, in fact, a strong deformation retraction (see in particular [58, Lemma 2.8(1)]). \square

Example 7.7. The skeleton of \mathbb{A}^1 is given by the half open line connecting 0 to ∞ , i.e., for trivially valued \mathbf{k} we have

$$\mathfrak{S}(X) = \left\{ |\cdot|_{0,r} \mid r \in [0, 1) \right\} \cup \left\{ |\cdot|_{\infty,r} \mid r \in [1, \infty) \right\}$$

in the notation of Example 7.1.

Let \mathbf{k} now be endowed with the trivial absolute value. As seen in [58, Sect. 2.1] the deformation retraction $\mathbf{p}_\Delta : X^{an} \rightarrow X^{an}$ restricts to a deformation retraction $\mathbf{p}_X : X^\natural \rightarrow X^\natural$, whose image is homeomorphic to the closure $\overline{\Delta}$ of Δ in $N_{\mathbb{R}}(\Delta)$. On T -invariant open affine subsets $X_\sigma = \text{Spec } \mathbf{k}[S_\sigma]$ this homeomorphism is induced by the tropicalization map

$$\begin{aligned} \text{trop}_X : X^\natural &\longrightarrow \text{Hom}(S_\sigma, \overline{\mathbb{R}}_{\geq 0}) \\ x &\longrightarrow (s \mapsto -\log |\chi^s|_x) \end{aligned}$$

into the extended cone $\bar{\sigma} = \text{Hom}(S_\sigma, \overline{\mathbb{R}}_{\geq 0})$.

7.3 The Case of Logarithmic Schemes

7.3.1 Zariski Logarithmic Schemes

Suppose that \mathbf{k} is endowed with the trivial absolute value and let X be logarithmically smooth over \mathbf{k} . In [58] Thuillier constructs a strong deformation retraction $\mathbf{p}_X : X^\natural \rightarrow X^\natural$ onto a closed subset $\mathfrak{S}(X)$ of X^\natural , the *skeleton* of X^\natural . We summarize the basic properties of this construction in the following Proposition 7.8. We refer to Definition 5.6 for the notion of small Zariski logarithmic schemes.

Proposition 7.8. (i) *The construction of \mathbf{p}_X is functorial with respect to logarithmic morphisms, i.e., given a logarithmic morphism $f : X \rightarrow Y$, there is a continuous map $f^\mathfrak{S} : \mathfrak{S}(X) \rightarrow \mathfrak{S}(Y)$ that makes the diagram*

$$\begin{array}{ccc}
 X^{\triangleright} & \xrightarrow{\mathbf{p}_X} & \mathfrak{S}(X) \\
 f^{\triangleright} \downarrow & & \downarrow f^{\mathfrak{S}} \\
 Y^{\triangleright} & \xrightarrow{\mathbf{p}_Y} & \mathfrak{S}(Y)
 \end{array}$$

commute. Moreover, if f is logarithmically smooth, then $f^{\mathfrak{S}}$ is the restriction of f^{\triangleright} to $\mathfrak{S}(X)$.

- (ii) For every strict étale neighborhood $U \rightarrow X$ that is small with respect to $x \in U$ and for every strict étale morphism $\gamma : U \rightarrow Z$ into a $\gamma(x)$ -small toric variety Z the analytic map γ^{\triangleright} induces a homeomorphism $\gamma^{\mathfrak{S}} : \mathfrak{S}(U) \xrightarrow{\sim} \mathfrak{S}(Z)$ that makes the diagram commute.

$$\begin{array}{ccc}
 U^{\triangleright} & \xrightarrow{\mathbf{p}_U} & \mathfrak{S}(U) \\
 \gamma^{\triangleright} \downarrow & & \sim \downarrow \gamma^{\mathfrak{S}} \\
 Z^{\triangleright} & \xrightarrow{\mathbf{p}_Z} & \mathfrak{S}(Z)
 \end{array}$$

- (iii) The skeleton $\mathfrak{S}(X)$ is the colimit of all skeletons $\mathfrak{S}(U)$ associated to strict étale morphisms $U \rightarrow X$ from a Zariski logarithmic scheme U that is small and the deformation retraction $\mathbf{p}_X : X^{\triangleright} \rightarrow X^{\triangleright}$ is induced by the universal property of colimits.

Suppose now that the logarithmic structure on X is defined in the Zariski topology. In this case, following [59], one can use the theory of Kato fans in order to define a tropicalization map $\text{trop}_X : X^{\triangleright} \rightarrow \bar{\Sigma}_X$ generalizing the map defined for toric varieties.

Let F be a fine and saturated Kato fan and consider the *cone complex*

$$\Sigma_F = F(\mathbb{R}_{\geq 0}) = \text{Hom}(\text{Spec } \mathbb{R}_{\geq 0}, F)$$

and the *extended cone complex*

$$\bar{\Sigma}_F = F(\bar{\mathbb{R}}_{\geq 0}) = \text{Hom}(\text{Spec } \bar{\mathbb{R}}_{\geq 0}, F) \supseteq \Sigma_F$$

associated to F . In order to describe the structure of $\bar{\Sigma}_F$ and Σ_F one can use the *structure map*

$$\begin{aligned}
 \rho : \bar{\Sigma}_F &\longrightarrow F \\
 u &\longmapsto u(\{\infty\})
 \end{aligned}$$

and the *reduction map*

$$r : \overline{\Sigma}_F \longrightarrow F$$

$$u \longmapsto u(\overline{\mathbb{R}}_{>0}) .$$

Proposition 7.9 ([61] Proposition 3.1). *The inverse image $r^{-1}(U)$ of an open affine subset $U = \text{Spec } P$ in $\overline{\Sigma}_F$ is the canonical compactification $\overline{\sigma}_U = \text{Hom}(P, \overline{\mathbb{R}}_{\geq 0})$ of a rational polyhedral cone $\sigma_U = \text{Hom}(P, \mathbb{R}_{\geq 0})$ and its relative interior $\overset{\circ}{\sigma}_U$ is given by $r^{-1}(x)$ for the unique closed point x in U .*

- (i) *If $V \subseteq U$ for open affine subsets $U, V \subseteq F$, then σ_V is a face of σ_U .*
- (ii) *For two open affine subsets U and V of F the intersection $\sigma_U \cap \sigma_V$ is a union of finitely many common faces.*

So the cone complex Σ_F is a rational polyhedral cone complex in the sense of [35]. It naturally carries the weak topology, in which a subset $A \subseteq \Sigma_F$ is closed if and only if the intersections $A \cap \sigma_U$ for all open affine subsets $U = \text{Spec } P$ of F are closed. The extended cone complex $\overline{\Sigma}_F$ is a canonical compactification of Σ_F , carrying the weak topology with the topology of pointwise convergence on $\sigma_U = \text{Hom}(P, \overline{\mathbb{R}}_{\geq 0})$ as local models.

- Proposition 7.10.** (i) *The reduction map $r : \overline{\Sigma}_F \rightarrow F$ is anti-continuous.*
 (ii) *The structure map $\rho : \overline{\Sigma}_F \rightarrow F$ is continuous.*
 (iii) *There is a natural stratification*

$$\bigsqcup_{x \in \text{Spec } F} \rho^{-1}(x) \simeq \overline{\Sigma}_F$$

of $\overline{\Sigma}_F$ by locally closed subsets.

Let X be a Zariski logarithmic scheme that is logarithmically smooth over \mathbf{k} and denote by $\phi_X : (X, \overline{\mathcal{O}}_X) \rightarrow F_X$ the *characteristic morphism* into its Kato fan F_X . We write Σ_X and $\overline{\Sigma}_X$ for the cone complex and the extended cone complex of F_X , respectively.

Following [59, Sect. 6.1] one can define the *tropicalization map* $\text{trop}_X \rightarrow \overline{\Sigma}_X$ as follows: a point $x \in X^{\mathbf{2}}$ can be represented by a morphism $\underline{x} : \text{Spec } R \rightarrow (X, \overline{\mathcal{O}}_X)$ for a valuation ring R extending \mathbf{k} . Its image $\text{trop}_X(x)$ in $\overline{\Sigma}_X = \text{Hom}(\text{Spec } \overline{\mathbb{R}}_{\geq 0}, F_X)$ is defined to be the composition

$$\text{Spec } \overline{\mathbb{R}}_{\geq 0} \xrightarrow{\text{val}^\#} \text{Spec } R \xrightarrow{\underline{x}} (X, \overline{\mathcal{O}}_X) \xrightarrow{\phi_X} F_X,$$

where $\text{val}^\#$ is the morphism $\text{Spec } \overline{\mathbb{R}}_{\geq 0} \rightarrow \text{Spec } R$ induced by the valuation $\text{val} : R \rightarrow \overline{\mathbb{R}}_{\geq 0}$ on R .

Proposition 7.11 ([59] Proposition 6.2).

(i) *The tropicalization map is well-defined and continuous. It makes the diagrams*

$$\begin{array}{ccc}
 X^\triangleright & \xrightarrow{\text{trop}_X} & \bar{\Sigma}_X \\
 r \downarrow & & r \downarrow \\
 (X, \bar{\mathcal{O}}_X) & \xrightarrow{\phi_X} & F_X
 \end{array}
 \qquad
 \begin{array}{ccc}
 X^\triangleright & \xrightarrow{\text{trop}_X} & \bar{\Sigma}_X \\
 \rho \downarrow & & \rho \downarrow \\
 (X, \bar{\mathcal{O}}_X) & \xrightarrow{\phi_X} & F_X
 \end{array}$$

commute.

(ii) *A morphism $f : X \rightarrow X'$ of Zariski logarithmic schemes, both logarithmically smooth and of finite type over \mathbf{k} , induces a continuous map $\bar{\Sigma}(f) : \bar{\Sigma}_X \rightarrow \bar{\Sigma}_{X'}$ such that the diagram commutes. The association $f \mapsto \bar{\Sigma}(f)$ is functorial in f .*

$$\begin{array}{ccc}
 X^\triangleright & \xrightarrow{\text{trop}_X} & \bar{\Sigma}_X \\
 f^\triangleright \downarrow & & \downarrow \bar{\Sigma}(f) \\
 (X')^\triangleright & \xrightarrow{\text{trop}_{X'}} & \bar{\Sigma}_{X'}
 \end{array}$$

Corollary 7.12 (Strata-Cone Correspondence, [61] Corollary 3.5). *There is an order-reversing one-to-one correspondence between the cones in Σ_X and the strata of X . Explicitly it is given by*

$$\mathring{\sigma} \mapsto r(\text{trop}_X^{-1}(\mathring{\sigma}))$$

for a relatively open cone $\mathring{\sigma} \subseteq \Sigma_X$ and

$$E \mapsto \text{trop}_X(r^{-1}(E) \cap X_0^{an})$$

for a stratum E of X .

Proof. This is an immediate consequence of the commutativity of

$$\begin{array}{ccc}
 X^\triangleright & \xrightarrow{\text{trop}_X} & \bar{\Sigma}_X \\
 r \downarrow & & r \downarrow \\
 X & \xrightarrow{\phi_X} & F_X
 \end{array}$$

from Propositions 7.11, 7.9, and the fact that ϕ_X sends every point in a stratum $E = E(\xi)$ to its generic point ξ . □

Corollary 7.13 ([61] Corollary 3.6). *The tropicalization map induces a continuous map $X^{\natural} \cap X_0^{an} \rightarrow \Sigma_F$.*

Proof. This follows from the commutativity of

$$\begin{array}{ccc} X^{\natural} & \xrightarrow{\text{trop}_X} & \bar{\Sigma}_X \\ \rho \downarrow & & \rho \downarrow \\ (X, \bar{\mathcal{O}}_X) & \xrightarrow{\phi_X} & F_X \end{array}$$

and the observations that $\rho^{-1}(X_0) = X^{\natural} \cap X^{an}$ as well as $\rho^{-1}(X_0 \cap F_X) = \Sigma_F$. \square

7.3.2 Étale Logarithmic Schemes

Let X be an étale logarithmic scheme that is logarithmically smooth over \mathbf{k} . We can define the *generalized extended cone complex* associated to X as the colimit of all $\bar{\Sigma}_{X'}$ taken over all strict étale morphisms $X' \rightarrow X$ from a Zariski logarithmic scheme X' . The tropicalization map $\text{trop}_X : X^{\natural} \rightarrow \bar{\Sigma}_X$ is induced by the universal property of colimits.

In analogy with Proposition 7.6 we have the following compatibility result stating that \mathbf{p}_X and trop_X are equal up to a natural homeomorphism.

Theorem 7.14 ([59] Theorem 1.2). *Suppose that X is logarithmically smooth over \mathbf{k} . There is a natural homeomorphism $J_X : \bar{\Sigma}_X \xrightarrow{\sim} \mathfrak{S}(X)$ making the diagram*

$$\begin{array}{ccc} & X^{\natural} & \\ \mathbf{p}_X \swarrow & & \searrow \text{trop}_X \\ \mathfrak{S}(X) & \xleftarrow[\sim]{J_X} & \bar{\Sigma}_X \end{array}$$

commute.

Moreover, we obtain the following Corollary of Proposition 7.11 (ii).

Corollary 7.15 ([59] Theorem 1.1). *A morphism $f : X \rightarrow X'$ of logarithmic schemes, logarithmically smooth and of finite type over \mathbf{k} , induces a continuous map $\bar{\Sigma}(f) : \bar{\Sigma}_X \rightarrow \bar{\Sigma}_{X'}$ such that the natural diagram*

$$\begin{array}{ccc}
 X^\natural & \xrightarrow{\text{trop}_X} & \bar{\Sigma}_X \\
 f^\natural \downarrow & & \downarrow \bar{\Sigma}(f) \\
 (X')^\natural & \xrightarrow{\text{trop}_{X'}} & \bar{\Sigma}_{X'}
 \end{array}$$

is commutative. The association $f \mapsto \bar{\Sigma}(f)$ is functorial in f .

8 Non-Archimedean Geometry of Artin Fans

8.1 Analytification of Artin Fans

In Sect. 7.3 we have seen that the extended cone complex $\bar{\Sigma}_F$ associated to a Kato fan F has topological properties analogous to the non-Archimedean analytic space X^\natural associated to a scheme X of finite type over \mathbf{k} . Moreover, if X is a Zariski logarithmic scheme that is logarithmically smooth over k , the tropicalization map $\text{trop}_X : X^\natural \rightarrow \bar{\Sigma}_X$ is the “analytification” of the characteristic morphism $\phi_X : (X, \mathcal{O}_X) \rightarrow F_X$. Using the theory of Artin fans we can make this analogy more precise, and even generalize the construction of trop_X to all logarithmic schemes.

Let \mathbf{k} be endowed with the trivial absolute value. As explained in [62, Sect. V.3] the $(\cdot)^\natural$ -functor, originally constructed in [58], generalizes to a pseudofunctor from the 2-category of algebraic stacks locally of finite type over \mathbf{k} into the category of non-Archimedean analytic stacks, such that whenever $[U/R] \simeq \mathcal{X}$ is a groupoid presentation of an algebraic stack \mathcal{X} we have a natural equivalence $[U^\natural/R^\natural] \simeq \mathcal{X}^\natural$. We refer the reader to [52, 60], and [64, Sect. 6] for background on the theory of non-Archimedean analytic stacks.

Let \mathcal{X} be an algebraic stack locally of finite type over \mathbf{k} . Then the underlying topological space $|\mathcal{X}^\natural|$ of the analytic stack \mathcal{X} can be identified with the set of equivalence classes of pairs (R, ϕ) consisting of a valuation ring R extending \mathbf{k} and a morphism $\phi : \text{Spec } R \rightarrow \mathcal{X}$. Two such pairs (ϕ, R) and (ψ, S) are said to be *equivalent*, if there is a valuation ring \mathcal{O} extending both R and S such that the diagram

$$\begin{array}{ccc}
 \text{Spec } \mathcal{O} & \longrightarrow & \text{Spec } S \\
 \downarrow & & \downarrow \psi \\
 \text{Spec } R & \xrightarrow{\phi} & X
 \end{array}$$

is 2-commutative. The topology on $|\mathcal{X}^\triangleright|$ is the coarsest making all maps $|U^\triangleright| \rightarrow |\mathcal{X}^\triangleright|$ induced by surjective flat morphisms $U \rightarrow \mathcal{X}$ from a scheme U locally of finite type over \mathbf{k} onto X into a topological quotient map.

Suppose that X is a logarithmic scheme that is logarithmically smooth and of finite type over \mathbf{k} . Consider the natural strict morphism $X \rightarrow \mathcal{A}_X$ into the Artin fan associated to X as constructed in Sect. 5.

Theorem 8.1 ([62]). *There is a natural homeomorphism*

$$\mu_X : |\mathcal{A}_X^\triangleright| \xrightarrow{\sim} \bar{\Sigma}_X$$

that makes the diagram

$$\begin{array}{ccc} & X^\triangleright & \\ \phi_X^\triangleright \swarrow & & \searrow \text{trop}_X \\ |\mathcal{A}_X^\triangleright| & \xrightarrow[\mu_X]{\sim} & \bar{\Sigma}_X \end{array}$$

commute.

So, by applying the functor $(\cdot)^\triangleright$ to the morphism $X \rightarrow \mathcal{A}_X$ we obtain the tropicalization map on the underlying topological spaces. Note that this construction also works for étale logarithmic schemes and we do not have to take colimits as in Sect. 7.3.2 (they are already taken in the construction of \mathcal{A}_X). Theorem 7.14 immediately yields the following Corollary.

Corollary 8.2. *There is a natural homeomorphism*

$$\tilde{\mu}_X : |\mathcal{A}_X^\triangleright| \xrightarrow{\sim} \mathfrak{S}(X)$$

that makes the diagram

$$\begin{array}{ccc} & X^\triangleright & \\ \phi_X^\triangleright \swarrow & & \searrow \mathbf{p}_X \\ |\mathcal{A}_X^\triangleright| & \xrightarrow[\tilde{\mu}_X]{\sim} & \mathfrak{S}(X) \end{array}$$

commute.

8.2 Stack Quotients and Tropicalization

Let X be a T -toric variety over \mathbf{k} . In this case, Theorem 8.1 precisely says that on the underlying topological spaces the tropicalization map $\text{trop}_X : X^{\mathfrak{Z}} \rightarrow \overline{\Sigma}_X$ is nothing but the analytic stack quotient map $X^{\mathfrak{Z}} \rightarrow [X^{\mathfrak{Z}}/T^{\mathfrak{Z}}]$. Due to the favorable algebro-geometric properties of toric varieties we can generalize this interpretation to general ground fields.

Let \mathbf{k} be any non-Archimedean field, and suppose that X is a T -toric variety defined by a rational polyhedral fan $\Delta \subseteq N_{\mathbb{R}}$ as in Sect. 7.2. Denote by T° the non-Archimedean analytic subgroup


$$\{x \in T^{an} \mid |\chi^m|_x = 1 \text{ for all } m \in M\}$$

of the analytic torus T^{an} . Note that, if \mathbf{k} is endowed with the trivial absolute value, then $T^{\circ} = T^{\mathfrak{Z}}$. In general, we can think of T° as a non-Archimedean analogue of the real n -torus $N_{S^1} = N \otimes_{\mathbb{Z}} S^1$ naturally sitting in $N_{\mathbb{C}^*} = N \otimes_{\mathbb{Z}} \mathbb{C}^*$. The T -operation on X induces an operation of T° on X^{an} .

Theorem 8.3 ([60] Theorem 1.1). *There is a natural homeomorphism*

$$\mu_{\Delta} : |[X^{an}/T^{\circ}]| \xrightarrow{\sim} N_{\mathbb{R}}(\Delta)$$

that makes the diagram



commute.

The proof Theorem 8.3 is based on establishing that the skeleton $\mathfrak{S}(X)$ of X is equal to the set of T° -invariant points of X^{an} . Then the statement follows, since by Ulirsch [60, Proposition 5.4 (ii)] the topological space $|[X^{an}/T^{\circ}]|$ is the colimit of the maps (Fig. 8)

$$T^{\circ} \times X^{an} \rightrightarrows X^{an} .$$




Fig. 8 Consider the affine line \mathbb{A}^1 over a trivially valued field k . The non-Archimedean unit circle \mathbb{G}_m° is given as the subset of elements in $x \in (\mathbb{A}^1)^{an}$ with $|t|_x = 1$, where t denotes a coordinate on \mathbb{A}^1 . The skeleton $\mathfrak{S}(\mathbb{A}^1)$ of $(\mathbb{A}^1)^{an}$ is the line connecting 0 to ∞ . It is precisely the set of “ \mathbb{G}_m° -invariant” points in $(\mathbb{A}^1)^{an}$, and therefore naturally homeomorphic to the topological space underlying $[(\mathbb{A}^1)^{an} / \mathbb{G}_m^\circ]$ (see [60, Example 6.2])

9 Where We Are, Where We Want to Go

9.1 Skeletons Fans and Tropicalization Over Non-Trivially Valued Fields

In almost all of our discussion, we have constructed Kato fans, Artin fans, and skeletons for a logarithmic variety over a trivially valued field k . One exception is the discussion of toric varieties in Sect. 7.2. This is quite useful and important: given a subvariety Y of a toric variety X , over a non-trivially valued field, the tropicalization $\text{trop}(Y)$ of Y is a polyhedral subcomplex of $N_{\mathbb{R}}(\Delta)$ which is not itself a fan. Much of the impact of tropical geometry relies on the way $\text{trop}(Y)$ reflects on the geometry of Y . So even though the toric variety X itself can be defined over a field with trivial valuation, the fact that Y —whose field of definition is non-trivially valued—has a combinatorial shadow is fundamental.

Can we define Artin fans over non-trivially valued fields? What should their structure be? In what generality can we canonically associate a skeleton with a logarithmic structure?

Some results in these directions are available in the recent paper of Gubler, Rabinoff, and Werner [27]. See also Werner’s contribution to this volume [63].

9.2 Improved Fans and Moduli Spaces

One of the primary applications of tropical geometry, and therefore of fans, is through moduli spaces. Mikhalkin’s correspondence theorem [41], Nishinou and Siebert’s vast extension [42], the work [15] of Cavalieri, Markwig, and Ranganathan, and the manuscript [6], all show that tropical moduli spaces $M^{\text{trop}}(X^{\text{trop}})$ of tropical curves in tropical varieties serve as good approximations of the tropicalization $\text{trop}(M(X))$ of moduli spaces $M(X)$ of algebraic curves in algebraic varieties. But the picture is not perfect:

1. Apart from very basic cases of the moduli space of curves itself [9] and genus-0 maps with toric targets [54] the tropical moduli space does not coincide with the tropicalization of the algebraic moduli space

$$M^{\text{trop}}(X^{\text{trop}}) \neq \text{trop}(M(X)).$$

2. Even when it does, it is always a coarse moduli space, lacking the full power of universal families.

These seem to be two distinct challenges, but experience shows that they are closely intertwined. It also seems that the following problem is part of the puzzle:

3. The construction to the Artin fan \mathcal{A}_X of a logarithmic scheme X is not functorial for all morphisms of logarithmic schemes.

The following program might inspire one to go some distance towards these challenges:

1. Good combinatorial moduli:
 - (a) Construct a satisfactory monoidal theory of cone spaces and cone stacks as described in Sects. 4.6.1 and 4.6.2.
 - (b) Construct a corresponding enhancement of $M_{g,n}^{\text{trop}}$ which is a fine moduli stack of tropical curves, with a universal family.

This is part of current work of Cavalieri, Chan, Ulirsch, and Wise.

2. Stacky Artin fans:
 - (a) Find a way to relax the representability condition in the definition of \mathcal{A}_X so that the enhanced moduli spaces above are canonically associated to the moduli stacks $M_{g,n}$
 - (b) Try to extend all of the above to moduli of maps.

Acknowledgements Research by Abramovich and Ulirsch is supported in part by NSF grant DMS-1162367 and BSF grant 2010255; Chen is supported in part by NSF grant DMS-1403271; Wise is supported by an NSA Young Investigator’s Grant, Award #H98230-14-1-0107.

References

1. Abramovich, D., Chen, Q.: Stable logarithmic maps to Deligne-Faltings pairs II. *Asian J. Math.* **18**(3), 465–488 (2014). MR 3257836
2. Abramovich, D., Fantechi, B.: Orbifold techniques in degeneration formulas. *Ann. Scuola Norm. Sup.* arXiv:1103.5132 (to appear)
3. Abramovich, D., Wise, J.: Invariance in logarithmic Gromov-Witten theory (2013). arXiv:1306.1222
4. Abramovich, D., Cadman, C., Wise, J.: Relative and orbifold Gromov-Witten invariants (2010). arXiv:1004.0981
5. Abramovich, D., Cadman, C., Fantechi, B., Wise, J.: Expanded degenerations and pairs. *Commun. Algebra* **41**(6), 2346–2386 (2013). MR 3225278
6. Abramovich, D., Chen, Q., Gross, M., Siebert, B.: Decomposition of degenerate Gromov-Witten invariants (2013, in preparation)
7. Abramovich, D., Marcus, S., Wise, J.: Comparison theorems for Gromov-Witten invariants of smooth pairs and of degenerations. *Ann. l’Institut Fourier* **64**(4), 1611–1667 (2014)
8. Abramovich, D., Chen, Q., Marcus, S., Wise, J.: Boundedness of the space of stable logarithmic maps. *J. Eur. Math. Soc.* arXiv:1408.0869 (to appear)
9. Abramovich, D., Caporaso, L., Payne, S.: The tropicalization of the moduli space of curves. *Annales de l’École Normale Supérieure* (4) **48**(4), 765–809 (2015)
10. Baker, M., Payne, S., Rabinoff, J.: Nonarchimedean geometry, tropicalization, and metrics on curves. *Algebr. Geom.* **3**(1), 63–105 (2016)
11. Berkovich, V.G.: *Spectral Theory and Analytic Geometry Over Non-Archimedean Fields*. *Mathematical Surveys and Monographs*, vol. 33. American Mathematical Society, Providence (1990). MR 1070709 (91k:32038)
12. Berkovich, V.G.: Étale cohomology for non-Archimedean analytic spaces. *Inst. Hautes Études Sci. Publ. Math.* **78**, 5–161 (1993/1994). MR 1259429 (95c:14017)
13. Berkovich, V.: Smooth p -adic analytic spaces are locally contractible. *Invent. Math.* **137**(1), 1–84 (1999). MR 1702143 (2000i:14028)
14. Cadman, C.: Using stacks to impose tangency conditions on curves. *Am. J. Math.* **129**(2), 405–427 (2007). MR MR2306040 (2008g:14016)
15. Cavaliere, R., Markwig, H., Ranganathan, D.: Tropicalizing the Space of Admissible Covers. *Math. Ann.* arXiv:1401.4626 (to appear)
16. Cavaliere, R., Marcus, S., Wise, J.: Polynomial families of tautological classes on $\mathcal{M}_{g,n}^{\text{tr}}$. *J. Pure Appl. Algebra* **216**(4), 950–981 (2012)
17. Chen, Q.: Stable logarithmic maps to Deligne-Faltings pairs I. *Ann. Math. (2)* **180**(2), 455–521 (2014). MR 3224717
18. Costello, K.: Higher genus Gromov-Witten invariants as genus zero invariants of symmetric products. *Ann. Math. (2)* **164**(2), 561–601 (2006)
19. Cox, D.A., Little, J.B., Schenck, H.K.: *Toric Varieties*. In: *Graduate Studies in Mathematics*, vol. 124. American Mathematical Society, Providence (2011). MR 2810322 (2012g:14094)
20. Danilov, V.I.: The geometry of toric varieties. *Uspekhi Mat. Nauk* **33**(2)(200), 85–134, 247 (1978). MR 495499 (80g:14001)
21. Demazure, M.: Sous-groupes algébriques de rang maximum du groupe de Cremona. *Ann. Sci. École Norm. Sup. (4)* **3**(4), 507–588 (1970). MR 0284446 (44 #1672)
22. Einsiedler, M., Kapranov, M., Lind, D.: Non-Archimedean amoebas and tropical varieties. *J. Reine Angew. Math.* **601**, 139–157 (2006). MR 2289207 (2007k:14038)
23. Fulton, W.: *Introduction to Toric Varieties*. *Annals of Mathematics Studies*, vol. 131. Princeton University Press, Princeton (1993); *The William H. Roever Lectures in Geometry*. MR 1234037 (94g:14028)
24. Gross, M., Siebert, B.: Logarithmic Gromov-Witten invariants. *J. Am. Math. Soc.* **26**(2), 451–510 (2013). MR 3011419

25. Gubler, W.: Tropical varieties for non-Archimedean analytic spaces. *Invent. Math.* **169**(2), 321–376 (2007). MR 2318559 (2008k:14085)
26. Gubler, W.: A guide to tropicalizations. In: *Algebraic and Combinatorial Aspects of Tropical Geometry*. Contemporary Mathematics, vol. 589, pp. 125–189. American Mathematical Society, Providence (2013). MR 3088913
27. Gubler, W., Rabinoff, J., Werner, A.: Skeletons and tropicalization. *Adv. Math.* (to appear)
28. Hrushovski, E., Loeser, F.: Non-archimedean Tame Topology and Stably Dominated Types. *Annals of Mathematics Studies*, vol. 192. Princeton University Press, Princeton (2016)
29. Ionel, E.-N.: GW invariants relative to normal crossing divisors. *Adv. Math.* **281**, 40–141 (2015)
30. Ionel, E.-N., Parker, T.H.: Relative Gromov-Witten invariants. *Ann. Math. (2)* **157**(1), 45–96 (2003). MR 1954264 (2004a:53112)
31. Ionel, E.-N., Parker, T.H.: The symplectic sum formula for Gromov-Witten invariants. *Ann. Math. (2)* **159**(3), 935–1025 (2004). MR 2113018 (2006b:53110)
32. Kajiwara, T.: Tropical toric geometry. In: *Toric Topology*. Contemporary Mathematics, vol. 460, pp. 197–207. American Mathematical Society, Providence (2008). MR 2428356 (2010c:14078)
33. Kato, K.: Logarithmic structures of Fontaine-Illusie. In: *Algebraic Analysis, Geometry, and Number Theory* (Baltimore, MD, 1988), pp. 191–224. Johns Hopkins University Press, Baltimore (1989). MR 1463703 (99b:14020)
34. Kato, K.: Toric singularities. *Am. J. Math.* **116**(5), 1073–1099 (1994). MR 1296725 (95g:14056)
35. Kempf, G., Knudsen, F.F., Mumford, D., Saint-Donat, B.: *Toroidal Embeddings. I*. Lecture Notes in Mathematics, vol. 339. Springer, Berlin (1973). MR 0335518 (49 #299)
36. Kim, B.: Logarithmic stable maps. In: *New Developments in Algebraic Geometry, Integrable Systems and Mirror Symmetry* (RIMS, Kyoto, 2008). Advanced Studies in Pure Mathematics, vol. 59, pp. 167–200. The Mathematical Society of Japan, Tokyo (2010). MR 2683209 (2011m:14019)
37. Li, J.: Stable morphisms to singular schemes and relative stable morphisms. *J. Differ. Geom.* **57**(3), 509–578 (2001). MR MR1882667 (2003d:14066)
38. Li, J.: A degeneration formula of GW-invariants. *J. Differ. Geom.* **60**(2), 199–293 (2002). MR MR1938113 (2004k:14096)
39. Li, A.-M., Ruan, Y.: Symplectic surgery and Gromov-Witten invariants of Calabi-Yau 3-folds. *Invent. Math.* **145**(1), 151–218 (2001). MR 1839289 (2002g:53158)
40. Macpherson, A.W.: *Skeleta in non-Archimedean and tropical geometry* (2013). arXiv:1311.0502
41. Mikhalkin, G.: Enumerative tropical algebraic geometry in \mathbb{R}^2 . *J. Am. Math. Soc.* **18**(2), 313–377 (2005). MR 2137980 (2006b:14097)
42. Nishinou, T., Siebert, B.: Toric degenerations of toric varieties and tropical curves. *Duke Math. J.* **135**(1), 1–51 (2006). MR 2259922 (2007h:14083)
43. Oda, T.: *Torus embeddings and applications*. Tata Institute of Fundamental Research Lectures on Mathematics and Physics, vol. 57. Tata Institute of Fundamental Research/Springer, Bombay/Berlin/New York (1978); Based on joint work with Katsuya Miyake. MR 546291 (81e:14001)
44. Oda, T.: *Convex bodies and algebraic geometry*. *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*, vol. 15. Springer, Berlin (1988); An introduction to the theory of toric varieties, Translated from the Japanese. MR 922894 (88m:14038)
45. Ogus, A.: *Lectures on logarithmic algebraic geometry* (2006)
46. Olsson, M.C.: Logarithmic geometry and algebraic stacks. *Ann. Sci. École Norm. Sup. (4)* **36**(5), 747–791 (2003)
47. Olsson, M.C.: The logarithmic cotangent complex. *Math. Ann.* **333**(4), 859–931 (2005). MR 2195148 (2006j:14017)
48. Parker, B.: *Gromov Witten invariants of exploded manifolds* (2011). arXiv:1102.0158

49. Parker, B.: Exploded manifolds. *Adv. Math.* **229**(6), 3256–3319 (2012). MR 2900440
50. Payne, S.: Analytification is the limit of all tropicalizations. *Math. Res. Lett.* **16**(3), 543–556 (2009). MR 2511632 (2010j:14104)
51. Popescu-Pampu, P., Stepanov, D.: Local tropicalization. In: *Algebraic and Combinatorial Aspects of Tropical Geometry*. Contemporary Mathematics, vol. 589. American Mathematical Society, Providence (2013)
52. Porta, M., Yu, T.Y.: Higher analytic stacks and GAGA theorems (2014). arXiv:1412.5166
53. Rabinoff, J.: Tropical analytic geometry, Newton polygons, and tropical intersections. *Adv. Math.* **229**(6), 3192–3255 (2012). MR 2900439
54. Ranganathan, D.: Moduli of rational curves in toric varieties and non-Archimedean geometry (2015). arXiv:1506.03754
55. Ranganathan, D.: Superabundant curves and the Artin fan. *IMRN* (to appear)
56. Tate, J.: Rigid analytic spaces. *Invent. Math.* **12**, 257–289 (1971). MR 0306196 (46 #5323)
57. Temkin, M.: Introduction to Berkovich analytic spaces. In: Ducros, A., Favre, C., Nicaise, J. (eds.) *Berkovich Spaces and Applications*. Lecture Notes in Mathematics, vol. 2119, pp. 3–66. Springer International Publishing, Cham (2015)
58. Thuillier, A.: Géométrie toroïdale et géométrie analytique non archimédienne. Application au type d’homotopie de certains schémas formels. *Manuscripta Math.* **123**(4), 381–451 (2007). MR 2320738 (2008g:14038)
59. Ulirsch, M.: Functorial tropicalization of logarithmic schemes: the case of constant coefficients (2013). arXiv:1310.6269
60. Ulirsch, M.: A geometric theory of non-Archimedean analytic stacks (2014). arXiv:1410.2216
61. Ulirsch, M.: Tropical compactification in log-regular varieties. *Math. Z.* **280**(1–2), 195–210 (2015). MR 3343903
62. Ulirsch, M.: Tropical geometry of logarithmic schemes, pp. viii+160. Ph.D. Thesis, Brown University, 2015
63. Werner, A.: Analytification and tropicalization over non-archimedean fields. In: *Proceedings of the 2013 and 2015 Simons Symposia on Nonarchimedean and Tropical Geometry* (2015, to appear). arXiv:1506.04846
64. Yu, T.Y.: Gromov compactness in non-archimedean analytic geometry. *J. Reine Angew. Math.* (to appear)

Introduction to Adic Tropicalization

Tyler Foster

Abstract This is an expository article on the adic tropicalization of algebraic varieties. We outline joint work with Sam Payne in which we put a topology and structure sheaf of local topological rings on the exploded tropicalization. The resulting object, which blends polyhedral data of the tropicalization with algebraic data of the associated initial degenerations, is called the *adic tropicalization*. It satisfies a theorem of the form “Huber analytification is the limit of all adic tropicalizations.” We explain this limit theorem in the present article, and illustrate connections between adic tropicalization and the curve complexes of O. Amini and M. Baker.

Keywords Non-archimedean geometry • Tropicalization • Huber adic spaces • limits of tropicalizations

1 Introduction

There is a close relationship between tropical geometry and the geometry of degenerations of algebraic varieties. This relationship takes a particularly suggestive form in much of the recent work that reinterprets mirror symmetry and enumerative geometry in terms of fibrations of complex varieties over affine manifolds. [8, 9, 16, 17]. An interesting manifestation of these ideas appears in the recent work of Parker [19, 21, 22]. Working in the paradigm of symplectic manifolds, with a view toward applications to pseudoholomorphic curve counting, Parker defines topological spaces called *exploded fibrations*, which he uses to construct log Gromov–Witten invariants [20], as required by the Gross–Siebert program. Roughly speaking, an exploded fibration is a topological space with structure semiring, which can be glued along boundary strata from pieces of the form $(\mathbb{C}^\times)^n \times P$, where P is an integral polytope in \mathbb{R}^m , for some m . For further details, see [19, Sect. 4] and

T. Foster (✉)

Department of Mathematics, University of Michigan, Ann Arbor, MI 48109, USA
e-mail: tyfoster@umich.edu

© Springer International Publishing Switzerland 2016
M. Baker, S. Payne (eds.), *Nonarchimedean and Tropical Geometry*,
Simons Symposia, DOI 10.1007/978-3-319-30945-3_10

337

[21, Sect. 3]. An illustrative example of an exploded fibration is the *exploded curve* given by the fibration $p : Y \longrightarrow B$ of topological spaces where $B \subset \mathbb{R}^2$ is the union of the subspaces

$$B_1 = \{(x, 0) \in \mathbb{R}^2 : x \geq 0\}, \quad B_2 = \{(0, y) \in \mathbb{R}^2 : y \geq 0\}, \quad \text{and} \\ B_3 = \{(x, x) \in \mathbb{R}^2 : x \leq 0\},$$

and where the space Y is glued from pieces

$$B_1 \times \mathbb{C}^\times - ((0, 0), 0), \quad B_2 \times \mathbb{C}^\times - ((0, 0), 1), \quad \text{and} \quad B_3 \times \mathbb{C}^\times - ((0, 0), \infty)$$

along an inclusion of $\{(0, 0)\} \times (\mathbb{P}^1_{\mathbb{C}} - \{0, 1, \infty\})$ into each $B_i \times \mathbb{C}^\times$, $i = 1, 2, 3$ (Fig. 1).


A tropical geometer recognizes the exploded curve above as a close analogue to the family of initial degenerations over all rational points in the tropicalization of an algebraic K -curve X . In more detail, let K be a non-Archimedean field, which is to say that K is complete with respect to a non-Archimedean absolute value $|\cdot| : K^\times \longrightarrow \mathbb{R}_{\geq 0}$. Assume that K is algebraically closed, with value group $\Gamma \stackrel{\text{def}}{=} -\log |K^\times| \subset \mathbb{R}$. Consider an n -dimensional split algebraic torus \mathbb{T} with character lattice M and cocharacter lattice $N = \text{Hom}_{\mathbb{Z}}(M, \mathbb{Z})$. Each vector $v \in N_{\mathbb{R}} \stackrel{\text{def}}{=} \mathbb{R} \otimes_{\mathbb{Z}} N$ determines the *tilted group ring* $R[U_v]$, the ring generated by all those monomials $a\chi^u$ for which $\langle u, v \rangle - \log |a| \geq 0$. Its spectrum $T_v \stackrel{\text{def}}{=} \text{Spec } R[U_v]$ is an R -scheme with generic fiber $T_{v,K} \cong \mathbb{T}$. When $v \in N_{\Gamma} \stackrel{\text{def}}{=} \Gamma \otimes_{\mathbb{Z}} N$, this ring $R[U_v]$ is finitely generated, with special fiber of the same dimension as its generic fiber \mathbb{T} . If X is any K -scheme, realized as a closed subvariety of the K -torus \mathbb{T} via some closed embedding $\iota : X \hookrightarrow \mathbb{T}$, then the *initial degeneration* of X at v , denoted $\text{in}_v X$, is the special fiber $(\overline{X})_k$ of the closure \overline{X} inside T_v .

The Fundamental Theorem of Tropical Geometry [18, Theorem 3.2.4] says that the tropicalization of X is the closure, in $N_{\mathbb{R}}$, of the set of those vectors $v \in N_{\Gamma}$ at which the initial degeneration $\text{in}_v X$ is nonempty:

$$\text{Trop}(X, \iota) = \overline{\{v \in N_{\Gamma} : \text{in}_v X \neq \emptyset\}}.$$

This means that we can retain more information about the subvariety $X \subset \mathbb{T}$ by remembering not just the points of the tropical variety $\text{Trop}(X, \iota)$, but the points of

Fig. 1 An example of an exploded curve in the sense of Parker [21]



all initial degenerations of X at all vectors $v \in N_\Gamma$. Payne makes use of this set-theoretical union

$$\bigsqcup_{v \in \text{Trop}(X) \cap N_\Gamma} |\text{in}_v X|$$

in order to study the fibers of the tropicalization map $X(K) \rightarrow \text{Trop}(X)$ [24, 25]. He refers to this set as the “exploded tropicalization,” taking inspiration from Parker’s exploded fibrations. In the present paper, we keep track of initial degenerations at all points of $\text{Trop}(X)$, Γ -rational or not, and define the *exploded tropicalization* to be the union

$$\mathfrak{Trop}(X) \stackrel{\text{def}}{=} \bigsqcup_{v \in \text{Trop}(X)} |\text{in}_v X|.$$

One major difference between Parker and Payne’s exploded objects is the presence of both a topology *and* a structure sheaf (of semirings) on Parker’s exploded fibrations, whereas Payne’s exploded tropicalizations are mere sets. The present paper is an exposition that complements the forthcoming research paper [5], joint with Sam Payne, in which we explain how to put both a topological structure and a structure sheaf (of topological rings) on the exploded tropicalization $\mathfrak{Trop}(X)$ of any closed subvariety X of a torus \mathbb{T} over a non-Archimedean field K . Our construction turns $\mathfrak{Trop}(X)$ into a locally topologically ringed space. We denote this space $\text{Ad}(X)$, and refer to it as the *adic tropicalization* of X in \mathbb{T} .

The locally topologically ringed space $\text{Ad}(X)$ interacts with Huber’s adic spaces [13, 14] in much the same way that tropical varieties interact with Berkovich spaces. In [23], Payne shows that if X is a quasi-projective variety over an algebraically closed non-Archimedean field K , then the topological space underlying the Berkovich analytification X^{an} is the inverse limit of all the tropical varieties $\text{Trop}(X, \iota)$, taken over the inverse system formed by all closed embeddings $\iota : X \hookrightarrow Y_\Sigma$ into quasi-projective toric varieties Y_Σ . In [6], Gross, Payne, and the author give a criterion extending this result to the case of certain non-quasi-projective varieties X over an algebraically closed, non-trivially valued non-Archimedean field K . In particular, [6, Theorem 1.2] X says that if X is a closed subvariety of a toric K -variety, then the topological space underlying X^{an} is the inverse limit of tropical varieties $\text{Trop}(X, \iota)$, taken over the inverse system formed by all closed embeddings $\iota : X \hookrightarrow Y_\Sigma$ into arbitrary toric varieties Y_Σ . One of the central results of the forthcoming paper [5] is a theorem stating that the Huber analytification X^{ad} of a closed subvariety X of a proper toric variety is isomorphic, as a locally topologically ringed space, to the inverse limit of the adic tropicalizations of X associated with all closed embeddings of X into proper toric varieties. Section 4 of the present paper provides a brief exposition on this result. See Theorem 4.4 and Proposition 4.8 below for precise statements.

Because the present paper is expository, it omits most proofs. The interested reader can find all the proofs in [5]. The present paper makes up for its lack of proofs by providing concrete details and examples that do not appear in [5]. In what follows, we focus more heavily on the basic geometry of adic spaces, which

underlies our geometrization of $\mathfrak{T}^{\text{top}}(X)$, and on concrete examples of the resulting topological spaces and their relationship to the metrized curve complexes of Amini and Baker [1].

2 Preliminaries on Huber Analytification

In the present Sect. 2, we review the basics of Huber’s theory of adic spaces and Huber’s analytification functor, which produces an adic space from any K -scheme X . The original sources for this material are [13] and [14]. We also strongly recommend Wedhorn’s notes [28] for a rather accessible introduction.

The reader already well versed in the theory of adic spaces can skip ahead to Sect. 3.

Notation. Fix a *non-Archimedean* field K , i.e., a field complete with respect to some non-trivial, non-Archimedean valuation $v : K \rightarrow \mathbb{R} \sqcup \{\infty\}$. Assume that K is algebraically closed. Let R denote the ring of integers in K , let \mathfrak{m} denote the unique maximal ideal in R , and let $k = R/\mathfrak{m}$ denote the residue field. Choose a real number $0 < \varepsilon < 1$ once and for all, and let $|\cdot|_v : K \rightarrow \mathbb{R}_{\geq 0}$ denote the norm $|a|_v = \varepsilon^{\text{val}(a)}$ associated with v .

Recall that the *Tate algebra*, denoted $K\langle T_1, \dots, T_n \rangle$, is the subalgebra of $K[[T_1, \dots, T_n]]$ consisting of those formal power series $f = \sum a_{i_1, \dots, i_n} T_1^{m_{i_1}} \cdots T_n^{m_{i_n}}$ for which $|a_{i_1, \dots, i_n}|_v \rightarrow 0$ as $i_1 + \dots + i_n \rightarrow \infty$. The Tate algebra comes with the *Gauss norm* $\|\cdot\|_G : K\langle T_1, \dots, T_n \rangle \rightarrow \mathbb{R}_{\geq 0}$, which takes $f \mapsto \|f\|_G \stackrel{\text{def}}{=} \max |a_{i_1, \dots, i_n}|_v$. The Gauss norm induces a *quotient norm* $\|\cdot\|_q$ on any quotient $\text{pr} : K\langle T_1, \dots, T_n \rangle \twoheadrightarrow K\langle T_1, \dots, T_n \rangle/\mathfrak{a}$, defined according to

$$\|f\|_q \stackrel{\text{def}}{=} \inf \{ \|g\|_G : g \in \text{pr}^{-1}(f) \}.$$

This quotient norm $\|\cdot\|_q$ gives $K\langle T_1, \dots, T_n \rangle/\mathfrak{a}$ the structure of a commutative K -Banach algebra. A *K -affinoid algebra* is any commutative Banach algebra A admitting at least one isomorphism of K -Banach algebras $A \cong K\langle T_1, \dots, T_n \rangle/\mathfrak{a}$. We use “ $\|\cdot\|_q$ ” to denote the norm on an arbitrary affinoid algebra A , even when there is no explicit mention of the presentation of A that gives rise to $\|\cdot\|_q$.

2.1 Adic Spectra and Adic Spaces

For an arbitrary totally ordered abelian group Γ' , written multiplicatively, let $\{0\} \sqcup \Gamma'$ denote the totally ordered abelian semigroup where $0 < \gamma$ for all $\gamma \in \Gamma'$, and where $0 \cdot \gamma = 0$ for all $\gamma \in \{0\} \sqcup \Gamma'$. A *continuous seminorm* $\|\cdot\|_x : A \rightarrow \{0\} \sqcup \Gamma'$ is any homomorphism of multiplicative semigroups satisfying

$$\|0\|_x = 0, \quad \|1\|_x = 1, \quad \text{and} \quad \|a + b\|_x \leq \max(\|a\|_x, \|b\|_x),$$

such that the set $\{a \in A : \|a\|_x < \gamma\}$ is open for each $\gamma \in \Gamma'$. Two continuous seminorms

$$\| - \|_x : A \longrightarrow \{0\} \sqcup \Gamma_1 \quad \text{and} \quad \| - \|_y : A \longrightarrow \{0\} \sqcup \Gamma_2$$

are *equivalent* if there exists an inclusion $\alpha : \{0\} \sqcup \Gamma_1 \hookrightarrow \{0\} \sqcup \Gamma_2$ of ordered abelian semigroups such that $\alpha \circ \| - \|_x = \| - \|_y$.

Let A be a K -affinoid algebra. The norm $\| - \|_q$ on A determines the *power bounded subring* $A^\circ \stackrel{\text{def}}{=} \{f \in A : \|f\|_q \leq 1\}$. The *adic spectrum* of the pair (A, A°) , denoted $\text{Spa}(A, A^\circ)$, is the set of all equivalence classes of continuous seminorms $\| - \|_x : A \longrightarrow \{0\} \sqcup \Gamma$ such that $\|a\|_x \leq 1$ for all $a \in A^\circ$. The topology on $\text{Spa}(A, A^\circ)$ is the topology generated by its *rational subsets*, that is, the coarsest topology containing the images of all inclusions

$$\text{Spa}\left(A\left\langle \frac{f_1, \dots, f_m}{g} \right\rangle, A\left\langle \frac{f_1, \dots, f_m}{g} \right\rangle^\circ\right) \hookrightarrow \text{Spa}(A, A^\circ) ,$$

where $A\left\langle \frac{f_1, \dots, f_m}{g} \right\rangle$ is any rational K -algebra determined by a set of elements $\{f_1, \dots, f_m, g\} \subset A$ generating the unit ideal in A (see [2, Sect.2.2.2.(ii)] and [28, Sect.8.1]).

If A is an affinoid algebra, with adic spectrum $X = \text{Spa}(A, A^\circ)$, and if U is a rational subset of X , let $\mathcal{O}_X(U)$ denote a choice of rational affinoid algebra associated with U as above. These rational algebras induce a *structure presheaf* \mathcal{O}_X on the adic spectrum X , which returns, at each open subset $U' \subset X$, the topological ring $\mathcal{O}_X(U')$ given by the inverse limit

$$\mathcal{O}_X(U') = \varprojlim_{U \subset U'} \mathcal{O}_X(U)$$

taken over all rational subsets U of U' in X . The theorem [13, Theorem 2.2] implies that the resulting structure presheaf \mathcal{O}_X is in fact a sheaf. Furthermore, we can form the stalk $\mathcal{O}_{X,x}$ at each point x in $\text{Spa}(A, A^\circ)$. This stalk is a local topological ring, and the seminorm $\| - \|_x$ induces a continuous seminorm on $\mathcal{O}_{X,x}$ [13, Proposition 1.6].

2.1.1 Adic Spaces

Let **LRS** denote the category of locally ringed spaces. By a *locally topologically ringed space*, we mean a locally ringed space (Y, \mathcal{O}_Y) such that \mathcal{O}_Y comes with the structure of a sheaf of topological rings. Let **LTRS** denote the category of locally topologically ringed spaces.

An *adic space* over K is a locally topologically ringed space (X, \mathcal{O}_X) that comes equipped with a continuous seminorm $\| - \|_x$ on the stalk $\mathcal{O}_{X,x}$ at each point $x \in X$, and which admits an atlas $\{U_i \hookrightarrow X\}$ in **LTRS**, whose charts U_i are adic spectra of K -affinoid algebras, such that the seminorm $\| - \|_x$ on each stalk $\mathcal{O}_{X,x} = \mathcal{O}_{U_i,x}$ coincides with the seminorm described in the previous paragraph. Note that this means, in particular, that the structure presheaf \mathcal{O}_X on any adic space must be a sheaf. A *morphism* of adic spaces is any morphism $\varphi : (X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y)$ of locally topologically ringed spaces such that, at each point $x \in X$, the induced morphism

$$\varphi^\# : \mathcal{O}_{Y,\varphi(x)} \rightarrow \mathcal{O}_{X,x}$$

on stalks satisfies $\| - \|_x \circ \varphi^\# = \| - \|_{\varphi(x)}$. We let **Adic** $_K$ denote the category of adic spaces over K .

2.1.2 The Sheaf of Power Bounded Sections

Every adic space (X, \mathcal{O}_X) comes with a second sheaf of topological rings, called the *sheaf of power bounded sections*. This is the subsheaf $\mathcal{O}_X^\circ \subset \mathcal{O}_X$ that, on each open subset $U \subset X$, returns those sections $f \in \mathcal{O}_X(U)$ for which $\|f\|_x \leq 1$ at every point $x \in U$.

The sheaf of power bounded sections is closely related to the theory of degenerations of K -varieties to the special fiber in $\text{Spec } R$, and it will play an important role in our formulation of the statement that ‘‘Huber analytification is the limit of all exploded tropicalizations’’ (see Theorem 4.4 below).

2.2 Huber Analytification

Fix a separated scheme X of finite type over K . It is a space defined locally by the vanishing of algebraic functions, and has no analytic structure. The Berkovich analytification X^{an} provides one way to realize an analytic structure on X . There is also an adic space X^{ad} that we can associate to X , called the Huber analytification of X . Intuitively, the Huber analytification of X is the pullback of X to the category of adic spaces over K along the morphism $\text{Spa}(K, R) \rightarrow \text{Spec } K$ of locally ringed spaces.

To make this precise, observe that there is a forgetful functor $U : \mathbf{Adic}_K \rightarrow \mathbf{LRS}$ that takes a given adic space Y and forgets both the seminorms on its stalks $\mathcal{O}_{Y,y}$ and the topological structure on the sheaf \mathcal{O}_Y . Let \mathbf{LRS}_K denote the category of locally ringed spaces over $\text{Spec } K$. There is a canonical isomorphism

$$U(\text{Spa}(K, R)) \xrightarrow{\sim} \text{Spec } K \tag{1}$$

of locally ringed spaces, which lets us interpret the forgetful functor U as a functor of the form

$$U : \mathbf{Adic}_K \longrightarrow \mathbf{LRS}_K.$$

Definition 2.1 (Huber Analytification). If X is a separated scheme of finite type over K , then its *Huber analytification*, denoted X^{ad} , is any adic space over $\text{Spa}(K, R)$ that comes equipped with a morphism $U(X^{\text{ad}}) \rightarrow X$ of locally ringed spaces over $\text{Spec } K$ satisfying the following universal property:

(H) If Y is any adic space over K that comes equipped with a map $\varphi : U(Y) \rightarrow X$ of locally ringed spaces over $\text{Spec } K$, then there exists a unique morphism $\tilde{\varphi} : Y \rightarrow X^{\text{ad}}$ of adic spaces over K for which the diagram

$$\begin{array}{ccc}
 U(Y) & & \\
 \downarrow & \searrow \varphi & \\
 U(\tilde{\varphi}) & & X \\
 \downarrow & \nearrow & \\
 U(X^{\text{ad}}) & &
 \end{array}$$

commutes.

Remark 2.2. The universal property **(H)** gives the Huber analytification the structure of a functor

$$(-)^{\text{ad}} : \mathbf{Sch}_K \longrightarrow \mathbf{Adic}_K.$$

For an explicit construction of the adic space X^{ad} as a fiber product of locally ringed spaces, see [28, Sect. 8.7 and Definition 8.63]. Most relevant to our purposes is the alternate description of X^{ad} as an inverse limit over admissible formal models of X^{an} , as established by van der Put and Schneider in [27]. To this end, we briefly review the theory of admissible formal models of X .

2.2.1 Admissible Formal Models

Recall that \mathfrak{m} denotes the maximal ideal in R , and that $k = R/\mathfrak{m}$ denotes our residue field. The theory of admissible formal models provides us with a tool for degenerating adic spaces over K to algebraic varieties over k . These degenerations are realized as certain formal R -schemes. Not all formal R -schemes will do. For instance, the category of formal R -schemes includes k -schemes as a full subcategory. A k -scheme is an example of a formal R -scheme that contains no information whatsoever over the generic point of $\text{Spec } R$. We want to excise these from our category of interest.

We say that a topological R -algebra A is *admissible* if there exists an isomorphism

$$A \cong R\langle t_1, \dots, t_n \rangle / \mathfrak{a}$$

of topological rings, where $R\langle t_1, \dots, t_n \rangle / \mathfrak{a}$ has its \mathfrak{m} -adic topology, such that:

- (1) the ideal $\mathfrak{a} \subset R\langle t_1, \dots, t_n \rangle$ is finitely generated and
- (2) the ring $R\langle t_1, \dots, t_n \rangle / \mathfrak{a}$ is free of \mathfrak{m} -torsion.

An *admissible formal R -scheme* is any formal scheme \mathfrak{X} over the formal spectrum $\mathrm{Spf}R$ of R (see [12, Chp. II, Sect. 9] for details) that admits an open covering $\{\mathfrak{U}_i \hookrightarrow \mathfrak{X}\}$ by formal spectra $\mathfrak{U}_i = \mathrm{Spf}A_i$, such that each A_i is an admissible R -algebra.

Each admissible formal R -scheme \mathfrak{X} has an associated adic space $(\mathfrak{X}^{\mathrm{ad}}, \mathcal{O}_{\mathfrak{X}^{\mathrm{ad}}})$. To construct $\mathfrak{X}^{\mathrm{ad}}$, choose a formal affine cover $\{\mathfrak{U}_i \hookrightarrow \mathfrak{X}\}$, say, with $\mathfrak{U}_i = \mathrm{Spf}A_i$, where A_i is an admissible R -algebra. Then $K \otimes_R A_i$ is a K -affinoid algebra, and we can define

$$\mathfrak{U}_i^{\mathrm{ad}} \stackrel{\mathrm{def}}{=} \mathrm{Spa}(K \otimes_R A_i, (K \otimes_R A_i)^\circ).$$

These adic spaces glue along the intersections $(\mathfrak{U}_i \cap \mathfrak{U}_j)^{\mathrm{ad}} = \mathfrak{U}_i^{\mathrm{ad}} \cap \mathfrak{U}_j^{\mathrm{ad}}$ to produce an adic space over K that we denote $\mathfrak{X}^{\mathrm{ad}}$. The adic space that we obtain in this way is independent of our choice of covering $\{\mathfrak{U}_i \hookrightarrow \mathfrak{X}\}$ (see [13, Proposition 4.1] for details).

Definition 2.3 (Admissible Formal Models of a Scheme). Let X be any scheme locally of finite type over K . An *admissible formal model* of X is a datum consisting of an admissible formal R -scheme \mathfrak{X} and an isomorphism $\mathfrak{X}^{\mathrm{ad}} \xrightarrow{\sim} X^{\mathrm{ad}}$ of adic spaces.

A *morphism* of admissible formal models of X is a morphism $\mathfrak{X}_1 \rightarrow \mathfrak{X}_2$ whose induced morphism $\mathfrak{X}_1^{\mathrm{ad}} \rightarrow \mathfrak{X}_2^{\mathrm{ad}}$ of K -analytic spaces commutes with the isomorphisms $\mathfrak{X}_1^{\mathrm{ad}} \xrightarrow{\sim} X^{\mathrm{ad}}$ and $\mathfrak{X}_2^{\mathrm{ad}} \xrightarrow{\sim} X^{\mathrm{ad}}$. We let \mathbf{AFS}_R denote the category of admissible formal R -schemes, and we let \mathbf{AFS}_R^X denote the category of admissible formal models of X .

2.2.2 Specializations of Adic Spaces

The adic space $\mathfrak{X}^{\mathrm{ad}}$ associated with the formal R -scheme \mathfrak{X} comes with a *specialization morphism*

$$\mathrm{sp}_{\mathfrak{X}} : (\mathfrak{X}^{\mathrm{ad}}, \mathcal{O}_{\mathfrak{X}^{\mathrm{ad}}}^\circ) \longrightarrow (\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}) \tag{2}$$

of locally topologically ringed spaces over $\mathrm{Spec} K$. Recall that $\mathcal{O}_{\mathfrak{X}^{\mathrm{ad}}}^\circ$ is the sheaf of power bounded sections on $\mathfrak{X}^{\mathrm{ad}}$. The fact that $\mathrm{sp}_{\mathfrak{X}}$ is a morphism of locally ringed

spaces is one feature of adic spaces that Berkovich spaces do *not* satisfy, as the reduction map $X^{\text{an}} \rightarrow \mathfrak{X}_k$ for Berkovich spaces fails to be continuous.

The specialization morphism (2) satisfies the following universal property:

(H_f) If Y is any adic space over K that comes equipped with a morphism $\varphi : (Y, \mathcal{O}_Y^\circ) \rightarrow (\mathfrak{X}, \mathcal{O}_{\mathfrak{X}})$ of locally topologically ringed spaces over $\text{Spec } K$, then there exists a unique morphism $\tilde{\varphi} : (Y, \mathcal{O}_Y) \rightarrow (\mathfrak{X}^{\text{ad}}, \mathcal{O}_{\mathfrak{X}^{\text{ad}}})$ of adic spaces over K for which the diagram

$$\begin{array}{ccc}
 (Y, \mathcal{O}_Y^\circ) & \xrightarrow{\varphi} & (\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}) \\
 \downarrow \tilde{\varphi} & & \uparrow \text{sp}_{\mathfrak{X}} \\
 (\mathfrak{X}^{\text{ad}}, \mathcal{O}_{\mathfrak{X}^{\text{ad}}}^\circ) & &
 \end{array}$$

commutes.

Remark 2.4. The universal property **(H_f)** gives us a functor $(-)^{\text{ad}} : \mathbf{AFS}_R \rightarrow \mathbf{Adic}_K$ into the category of adic spaces over K . If X is a separated finite-type K -scheme, then the resulting specialization morphisms

$$\text{sp}_{\mathfrak{X}} : (X^{\text{ad}}, \mathcal{O}_{X^{\text{ad}}}^\circ) \rightarrow (\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}), \tag{3}$$

induced by the specialization morphisms (2) commute with morphisms in \mathbf{AFS}_R^X .

Gillam proved in [7, Corollary 5] that **LRS** contains all inverse limits. It is not difficult to show that one can compute any inverse limit $\varprojlim (Y_i, \mathcal{O}_{Y_i})$ in **LTRS** by first computing the inverse limit (Y, \mathcal{O}_Y) in **LRS** of the underlying locally ringed spaces, and then equipping each ring $\mathcal{O}_Y(U)$ with the finest topology for which each of the morphisms

$$\mathcal{O}_{Y_i}(V_i) \rightarrow \mathcal{O}_X(U), \tag{4}$$

for each open subset $V_i \subset Y_i$ mapping to U under $Y_i \rightarrow Y$, becomes continuous.

This means, in particular, that the inverse limit of an inverse system of formal schemes remains a topologically ringed space on which stalks are local rings. Thus the morphisms (3) induce a morphism

$$\text{sp} : (X^{\text{ad}}, \mathcal{O}_{X^{\text{ad}}}^\circ) \rightarrow \varprojlim_{\mathbf{AFS}_R^X} (\mathfrak{X}, \mathcal{O}_{\mathfrak{X}}) \tag{5}$$

of locally topologically ringed spaces.

Proposition 2.5. *The morphism (5) is an isomorphism of locally topologically ringed spaces.*

Proof. For a sketch of the proof of Proposition 2.5, see [26, Theorem 2.22]. □

3 The Exploded Tropicalization as a Locally Ringed Space

3.1 The Combinatorics of Toric Degenerations

In the present Sect. 3.1, we briefly review the combinatorial and algebraic geometry of toric degenerations of toric varieties. We work exclusively over the toric k -variety \mathbb{A}_k^1 , where $k = R/\mathfrak{m}$ is our original residue field. The reader already familiar with the geometry of toric degenerations over \mathbb{A}_k^1 , as developed by Kempf et al. in [15], may proceed to Sect. 3.2, where we review the formalism of toric R -schemes introduced by Gubler in [10]. The theory of toric R -schemes is further developed by Gubler and Soto in [11].

Fix a split, n -dimensional algebraic torus \mathbb{T} over k , with character lattice M . We have a canonical isomorphism $\mathbb{T} \cong \text{Spec } k[M]$. Let $N \stackrel{\text{def}}{=} \text{Hom}_{\mathbb{Z}}(M, \mathbb{Z})$ be the cocharacter lattice, and let

$$\langle -, - \rangle : M \times N_{\mathbb{R}} \longrightarrow \mathbb{R}$$

denote the canonical pairing.

If Y_{Σ} is a proper toric variety over K , built from a complete n -dimensional fan Σ in $N_{\mathbb{R}}$, then we can freely produce degenerations of Y_{Σ} by writing down certain $(n+1)$ -dimensional fans extending Σ . Specifically, consider the $(n-1)$ -dimensional half-space $N_{\mathbb{R}} \times \mathbb{R}_{\geq 0}$, along with its projection to the second factor

$$\text{pr}_2 : N_{\mathbb{R}} \times \mathbb{R}_{\geq 0} \longrightarrow \mathbb{R}_{\geq 0}. \tag{6}$$

If Δ is any fan in $N_{\mathbb{R}} \times \mathbb{R}$ with support $N_{\mathbb{R}} \times \mathbb{R}_{\geq 0}$, such that

$$\Delta \cap (N_{\mathbb{R}} \times \{0\}) \cong \Sigma,$$

then the projection (6) induces a morphism of fans

$$\text{pr}_2 : \Delta \longrightarrow \{0, \mathbb{R}_{\geq 0}\} \quad \text{with} \quad \text{pr}_2^{-1}(0) = \Sigma. \tag{7}$$

This induces a torus-equivariant morphism of toric varieties

$$f_{\text{pr}_2} : Y_{\Delta} \longrightarrow \mathbb{A}^1 = \text{Spec } k[t]. \tag{8}$$

See Fig. 2 below for an example. Note that the fiber $\text{pr}_2^{-1}(1) = \Delta \cap (N_{\mathbb{R}} \times \{1\})$ is not a fan but a polyhedral complex. We identify it with a polyhedral complex C_{Δ} in $N_{\mathbb{R}}$ whose support is all of $N_{\mathbb{R}}$. See Fig. 3 below for an example.

Proposition 3.1. *The map $f_{\text{pr}_2} : Y_{\Delta} \longrightarrow \mathbb{A}^1$ induced by the projection (7) realizes the complex toric variety Y_{Δ} as an \mathbb{A}^1 -model of the toric variety glued from Σ . More precisely:*

- (i) The generic fiber of f_{pr_2} is the toric $k(t)$ -variety Y_Σ .
- (ii) The special fiber of f_{pr_2} at $t = 0$ is a union of toric varieties $Y_{\text{star}(v)}$ indexed by the vertices v of C_Δ , glued along torus-invariant divisors according to the higher-dimensional face incidences appearing in C_Δ .

Proof. The complex toric variety Y_Δ is glued from the affine open toric varieties $U_\delta \stackrel{\text{def}}{=} \text{Spec } k[S_\delta]$, where δ runs over all cones in Δ . Because Δ is a fan in $N_\mathbb{R} \times \mathbb{R}_{\geq 0}$, the semigroup S_δ associated with δ takes the form

$$S_\delta \stackrel{\text{def}}{=} \{(u, n) \in M \times \mathbb{Z} : \langle u, v \rangle + rn \geq 0 \text{ for all } (v, r) \in \delta\}. \tag{9}$$

Fix a monomial $a t^n \chi^u \in k[S_\delta]$, where $a \in k$ and $u \in M$. If $(u, 0) \notin S_{\delta \cap (N_\mathbb{R} \times \{0\})}$, then there exist some $(v', 0) \in \delta \cap (N_\mathbb{R} \times \{0\})$ such that $\langle u, v' \rangle < 0$. Thus, for some choice of $r \gg 0$, we have $\langle u, v + rv' \rangle < 0$. Because δ is closed under translation by $\delta \cap (N_\mathbb{R} \times \{0\})$, this implies that $u \in S_{\delta \cap (N_\mathbb{R} \times \{0\})}$. This proves part (i) of the proposition.

Part (ii) is a consequence of the orbit-cone correspondence for toric varieties [4, Sect. 3.2], upon observing that the special divisors in Y_Δ are exactly those codimension-1 orbit closures in Y_Δ that correspond to rays in Δ not lying in $N_\mathbb{R} \times \{0\}$. \square

Example 3.2. Let $M = \mathbb{Z}^2$, and let Σ be the complete, 2-dimensional fan in $N_\mathbb{R} = \mathbb{R}^2$ pictured at left in Fig. 3 below. Then $Y_\Sigma = \mathbb{P}^2_\mathbb{C}$. Let Δ be the 3-dimensional, \mathbb{Z} -admissible fan in $\mathbb{R}^2 \times \mathbb{R}_{\geq 0}$ pictured in Fig. 2 below.

Because Δ satisfies $\Delta \cap (\mathbb{R}^2 \times \{0\}) = \Sigma$, it produces a complex toric variety Y_Σ that comes equipped with a torus-equivariant map $Y_\Sigma \longrightarrow \mathbb{A}^1$. The generic fiber of the family is \mathbb{P}^2 over $k(t)$. The special fiber over $t = 0$ is a union of four complex surfaces: a copy of \mathbb{P}^2 , a copy of $\mathbb{P}^1 \times \mathbb{P}^1$, and two Hirzebruch surfaces. These four surfaces are in one-to-one correspondence with the vertices of the polyhedral decomposition $\Delta \cap (\mathbb{R}^2 \times \{1\})$ pictured at left in Fig. 3. More precisely, they are the toric varieties described by the stars at each vertex, and they are glued to one another along the torus-invariant divisors corresponding to the edges connecting their corresponding vertices in C_Δ .




Fig. 2 A map of fans $\text{pr}_2 : \Delta \longrightarrow \{0, \mathbb{R}_{\geq 0}\}$ inducing toric morphism of toric varieties $f_{\text{pr}_2} : Y_\Delta \longrightarrow \mathbb{A}^1$




Fig. 3 The fan $\Sigma \cong \Delta \cap (\mathbb{R}^2 \times \{0\})$ and the polyhedral complex $\Delta \cap (\mathbb{R}^2 \times \{1\})$

3.2 Algebraic Gubler Models

The fundamental observation underlying Gubler’s approach to integral toric geometry, already anticipated by Kempf et al. in [15], is that a significant part of the geometry of toric degenerations discussed in the previous Sect. 3.1 works just as well upon replacing the curve $\mathbb{A}^1 = \text{Spec } k[t]$ with the spectrum of the ring of integers R in any non-Archimedean field K . This requires that we replace the factor \mathbb{Z} in $M \times \mathbb{Z}$ with the value group $\Gamma = \text{val}(K^\times) \subset \mathbb{R}$ of K . The fact that we can express each semigroup S_δ appearing in the proof of Proposition 3.1 in the form (9) implies that we can express the corresponding semigroup rings as

$$k[S_\delta] = \left\{ \sum_{u \in M} a_u(t) \chi^u \in k[t, t^{-1}][M] : \langle u, v \rangle + r \text{ord}_t(a_u(t)) \geq 0 \text{ for all } (v, r) \in \delta \right\}, \tag{10}$$

where the coefficients $a_u(t)$ are Laurent polynomials in t , and where

$$\text{ord}_t : k(t)^\times \longrightarrow \mathbb{Z}$$

is the valuation on $k(t)$ given by order-of-vanishing at $t = 0$. In Gubler’s formalism, one replaces the ring $k[t, t^{-1}]$ in (10) with a non-Archimedean field K . This produces R -algebras that glue to give an R -model of the toric K -variety Y_Σ . The present Sect. 3.2 is devoted to a review of Gubler’s formalism, with slight upgrades to the setting of formal R -schemes. In the next Sect. 3.3, we use the resulting formal R -schemes to put a topology and locally ringed structure sheaf on each exploded tropicalization.

Definition 3.3. A Γ -admissible cone is any convex cone $\delta \subset N_\mathbb{R} \times \mathbb{R}_{\geq 0}$ that can be written as an intersection of the form

$$\delta = \bigcap_{i=1}^n \{(v, c) \in N \times \mathbb{R}_{\geq 0} : \langle u_i, v \rangle + \gamma_i c \geq 0\}$$

for vectors $(u_1, \gamma_1), \dots, (u_n, \gamma_n)$ in $M \times \Gamma$. A Γ -admissible fan is any fan Δ in $N_{\mathbb{R}} \times \mathbb{R}_{\geq 0}$ whose every cone δ is Γ -admissible.

Definition 3.4. If δ is a Γ -admissible cone in $N_{\mathbb{R}} \times \mathbb{R}_{\geq 0}$, its associated δ -tilted algebra is the ring

$$R[U_{\delta}] \stackrel{\text{def}}{=} \left\{ \sum_{u \in M} a_u \chi^u \in K[M] : \langle u, v \rangle + \text{val}(a_u) \cdot c \geq 0 \text{ for all } (v, c) \in \delta \right\}.$$

By [10, Proposition 11.3], every δ -tilted algebra $R[U_{\delta}]$ is in fact a finite-type, flat R -algebra. We let U_{δ} denote the R -scheme $U_{\delta} = \text{Spec } R[U_{\delta}]$.

Each inclusion $\delta' \hookrightarrow \delta$ of Γ -admissible cones in $N_{\mathbb{R}} \times \mathbb{R}_{\geq 0}$ realizing δ' as a face of δ induces a functorial, open embedding of R -schemes

$$U_{\delta'} \hookrightarrow U_{\delta}. \tag{11}$$

If Δ is a Γ -admissible fan in $N_{\mathbb{R}} \times \mathbb{R}_{\geq 0}$, then we can glue the affine R -schemes $\{U_{\delta}\}_{\delta \in \Delta}$ along these induced open embeddings (11) to obtain a single R -scheme Y_{Δ} . The R -scheme Y_{Δ} is locally of finite type over R . Gubler shows [10] that when the support of Δ is $N_{\mathbb{R}} \times \mathbb{R}_{\geq 0}$, the scheme Y_{Δ} is in fact proper over $\text{Spec } R$, with generic fiber canonically isomorphic to the toric K -variety Y_{Σ} glued from the fan

$$\Sigma \stackrel{\text{def}}{=} \Delta \cap (N_{\mathbb{R}} \times \{0\}),$$

upon interpreting Σ as a fan in $N_{\mathbb{R}}$. In short:

$$(Y_{\Delta})_K \cong Y_{\Sigma}. \tag{12}$$

Identifying $N_{\mathbb{R}} \times \{0\}$ with $N_{\mathbb{R}}$, we refer to the intersection $\Delta \cap (N_{\mathbb{R}} \times \{0\})$ as the *recession fan of Δ in $N_{\mathbb{R}}$* , and denote it $\text{rec}(\Delta)$. Just as in Proposition 3.1. (ii), the special fiber of Y_{Δ} over the unique maximal ideal in $\text{Spec } R$ is a union of toric k -varieties indexed by the vertices in the *height-1 polyhedral complex*

$$\Delta \cap (N_{\mathbb{R}} \times \{1\}), \tag{13}$$

and glued along torus-invariant divisors according to the incidence profile of higher-dimensional polyhedra in this height-1 complex.

We refer to Y_{Δ} as the (*algebraic*) *Gubler model* associated with the Γ -admissible fan Δ .

Example 3.5. Let $K = k((t))$, so that $R = k[[t]]$ and $\Gamma = \mathbb{Z}$. Then the fan Δ from Example 3.2 above, being integral, is Γ -admissible. If \mathbb{Y}_{Δ} denotes the toric

k -variety that we constructed in Example 3.2, equipped with its torus-equivariant fibration $f_{\text{pr}_2} : \mathbb{Y}_\Delta \rightarrow \mathbb{A}^1$, then the algebraic Gubler model associated with Δ is the fiber

$$Y_\Delta \cong \text{Spec } k[[t]] \times_{\mathbb{A}^1} \mathbb{Y}_\Delta$$

over the $k[[t]]$ -valued point $\text{Spec } k[[t]] \rightarrow \mathbb{A}^1$ supported at $t = 0$.

Example 3.6. Let $K = \mathbb{Q}_p$, so that $R = \mathbb{Z}_p$ and $\Gamma = \mathbb{Z}$. Define $M = \mathbb{Z}$ and $N = \mathbb{Z}$, so that the standard pairing is just multiplication. Let Δ be the \mathbb{Z} -admissible fan in $N_{\mathbb{R}} \times \mathbb{R}_{\geq 0} = \mathbb{R} \times \mathbb{R}_{\geq 0}$ pictured at left in Fig. 4 below.

Going from left to right, the 2-dimensional cones in Δ have corresponding tilted algebras

$$\mathbb{Z}_p[t^{-1}], \quad \mathbb{Z}_p[t, pt^{-1}] \cong \mathbb{Z}_p[x, y]/(xy - p), \quad \text{and} \quad \mathbb{Z}_p[p^{-1}t].$$

Thus Y_Δ describes the projective line over \mathbb{Q}_p degenerating to a pair of projective lines over \mathbb{F}_p intersecting at a node.

3.3 Formal Gubler Models

If Δ is a Γ -admissible fan in $N_{\mathbb{R}} \times \mathbb{R}_{\geq 0}$, whose support is all of $N_{\mathbb{R}} \times \mathbb{R}_{\geq 0}$, then we let \mathfrak{Y}_Δ denote the formal R -scheme obtained as the completion of Y_Δ along its special fiber,

$$\mathfrak{Y}_\Delta \stackrel{\text{def}}{=} \widehat{Y_\Delta}.$$

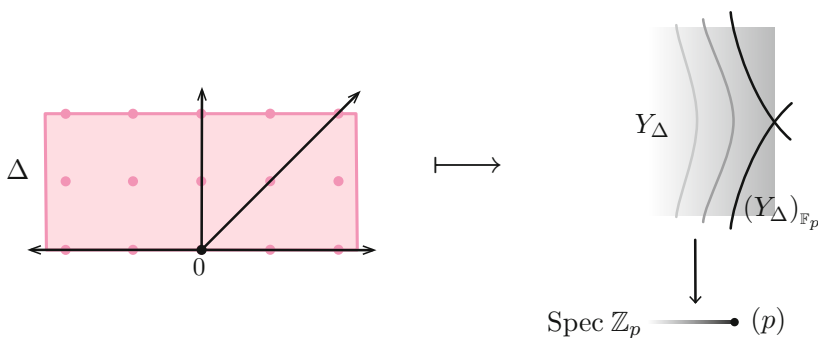


Fig. 4 A \mathbb{Z} -admissible fan in $\mathbb{R} \times \mathbb{R}_{\geq 0}$, and its associated \mathbb{Z}_p -model of $\mathbb{P}^1_{\mathbb{Q}_p}$

If, at each cone δ in Δ , we let $R[\mathfrak{U}_\delta]$ denote the formal completion

$$R[\mathfrak{U}_\delta] \stackrel{\text{def}}{=} \widehat{R[U_\delta]}$$

along the ideal generated by \mathfrak{m} in $R[U_\delta]$, then \mathfrak{Y}_Δ is glued from the formal spectra

$$\mathfrak{U}_\delta = \text{Spf } R[\mathfrak{U}_\delta].$$

Because each \mathfrak{U}_δ is an admissible formal R -scheme, this implies that \mathfrak{Y}_Δ is itself an admissible formal R -scheme.

Because \mathfrak{Y}_Δ is a formal completion, the “generic fiber” of \mathfrak{Y}_Δ is no longer a K -scheme as in the previous Sect. 3.2. Instead, [3, Sect. 2.4] tells us that in the formal setting, the topological K -algebras

$$K \otimes_R R[\mathfrak{U}_\delta], \text{ for } \delta \text{ in } \Delta,$$

are K -affinoid algebras. Their adic spectra

$$\mathfrak{U}_\delta^{\text{ad}} \stackrel{\text{def}}{=} \text{Spa}(K \otimes_R R[\mathfrak{U}_\delta], R[\mathfrak{U}_\delta]) \tag{14}$$

glue to produce an *adic generic fiber* $\mathfrak{Y}_\Delta^{\text{ad}}$ of the formal R -scheme \mathfrak{Y}_Δ . In place of the isomorphism of K -varieties (12), we have a canonical isomorphism of adic spaces

$$\mathfrak{Y}_\Delta^{\text{ad}} \cong Y_\Sigma^{\text{ad}}.$$

In the language of Definition 2.3, the formal R -scheme \mathfrak{Y}_Δ is an admissible formal model of the K -scheme Y_Σ .

3.4 Adic Tropicalization of a Closed Embedding

Fix a proper algebraic variety X over K . Even if X fails to be projective, we can still ask for closed embeddings into proper toric K -varieties. Let us suppose that X is a proper K -variety admitting at least one closed embedding into a proper toric K -variety

$$\iota : X \hookrightarrow Y_\Sigma,$$

where Σ is a complete fan in $N_\mathbb{R}$. Using the machinery of Sect. 3.2, we can produce a flat, proper, algebraic R -model Y_Δ of Y_Σ by making a choice of Γ -admissible fan Δ in $N_\mathbb{R} \times \mathbb{R}_{\geq 0}$ with support equal to $N_\mathbb{R} \times \mathbb{R}_{\geq 0}$. Our variety X sits in the generic

fiber of this model, and Gubler shows [10] that the closure \overline{X} of X in Y_Δ is itself a flat, proper R -model of X . Let

$$\mathfrak{X}_{(\Delta, \iota)} \stackrel{\text{def}}{=} \widehat{\overline{X}}$$

denote the formal completion of the algebraic R -model \overline{X} along its special fiber. This formal completion admits a canonical isomorphism $\mathfrak{X}_{(\Delta, \iota)}^{\text{ad}} \cong X^{\text{ad}}$.

Definition 3.7. The category of fans recessed over ι , denoted $\mathbf{REC}_{\Sigma, \iota}$, is the category with:

objects: Any locally finite Γ -admissible fan Δ in $N_{\mathbb{R}} \times \mathbb{R}_{\geq 0}$ with recession fan $\text{rec}(\Delta) = \Sigma$ and

morphisms: Any pair of locally finite Γ -admissible fans Δ' and Δ in $N_{\mathbb{R}} \times \mathbb{R}_{\geq 0}$ such that each cone δ' in Δ' is contained in at least one cone δ in Δ .

We refer to a morphism $\rho : \Delta' \rightarrow \Delta$ in $\mathbf{REC}_{\Sigma, \iota}$ as a *refinement* of Δ .

If $\rho : \Delta' \rightarrow \Delta$ is a refinement of Δ , then it induces a morphism $f_\rho : Y_{\Delta'} \rightarrow Y_\Delta$ of R -models of Y_Σ . This map of R -models itself restricts to a morphism

$$f_\rho : \mathfrak{X}_{(\Delta', \iota)} \rightarrow \mathfrak{X}_{(\Delta, \iota)}$$

of admissible formal R -models of X . In this way, the assignment $\Delta \mapsto \mathfrak{X}_{(\Delta, \iota)}$ becomes a functor $\mathbf{REC}_{\Sigma, \iota} \rightarrow \mathbf{AFS}_R^X$ that takes values in the category of admissible formal models of X^{an} .

Definition 3.8 (Adic Tropicalization). The *adic tropicalization* of the closed embedding $\iota : X \hookrightarrow Y_\Sigma$ is the inverse limit

$$(\text{Ad}(X, \iota), \mathcal{O}_{\text{Ad}(X, \iota)}) \stackrel{\text{def}}{=} \varprojlim_{\mathbf{REC}_{\Sigma, \iota}} (\mathfrak{X}_{(\Delta, \iota)}, \mathcal{O}_{\mathfrak{X}_{(\Delta, \iota)}})$$

in the category of locally topologically ringed spaces.

The utility of Definition 3.8 comes from the fact that, as an immediate consequence of Remark 2.4, it gives $\text{Ad}(X, \iota)$ the structure of a locally topologically ringed topological space. But Definition 3.8 obscures the relationship between $\text{Ad}(X, \iota)$ and the exploded tropicalization of X . To clarify this relationship, first consider the special case where $\iota : X' \hookrightarrow \mathbb{T}'$ is a closed embedding into a torus, with N' the cocharacter lattice of \mathbb{T}' . In this case, the *exploded tropicalization* is as described in Sect. 1: it is the union

$$\mathfrak{Trop}(X', \iota) \stackrel{\text{def}}{=} \bigsqcup_{v \in N'_{\mathbb{R}}} |\text{in}_v X'|.$$




Fig. 5 The ray τ_v spanned by a vector $(v, 1)$ in $N_{\mathbb{R}} \times \{1\}$, at *left*, and a Γ -admissible cone δ containing τ_v when τ_v is not itself Γ -rational

For a closed embedding $\iota : X \hookrightarrow Y_{\Sigma}$ into an arbitrary toric variety Y_{Σ} , the *extended exploded tropicalization*, denoted $\mathfrak{Top}(X, \iota)$, is the set-theoretical union

$$\mathfrak{Top}(X, \iota) = \bigsqcup_{\sigma \in \Sigma} \mathfrak{Top}(\iota^{-1}(O(\sigma)), \iota),$$

where $\iota^{-1}(O(\sigma))$ is the part of X that maps to the torus orbit $O(\sigma) \cong \mathbb{T}_{N/\sigma \cap N}$ under ι (see [4, Sect. 3.2]). In other words, the extended exploded tropicalization is the union of all exploded tropicalizations encoded in ι as we run over all locally closed torus-orbit strata in Y_{Σ} .

Note that each Γ -rational vector v in $N_{\mathbb{R}}$ determines a Γ -admissible ray $\tau_v \stackrel{\text{def}}{=} \mathbb{R}_{\ge 0}(v, 1)$ in $N_{\mathbb{R}} \times \mathbb{R}_{\ge 0}$, and vice versa, as depicted at left in Fig. 5. When v is *not* Γ -rational, the ray τ_v is no longer Γ -admissible, but there are still many Γ -admissible cones δ containing τ_v , as depicted at right in Fig. 5.

Proposition 3.9. *For any vector v in $N_{\mathbb{R}}$, we have a natural isomorphism*

$$\text{in}_v X \cong \varprojlim_{\delta \supset \tau_v} (\mathfrak{X}_{\delta})_k, \tag{15}$$

where the limit runs over all Γ -admissible cones δ in $N_{\mathbb{R}} \times \mathbb{R}_{\ge 0}$ that contain the ray τ_v . When v is Γ -rational, the above isomorphism (15) simplifies to an isomorphism

$$\text{in}_v X \cong (\mathfrak{X}_{\tau_v})_k.$$

One uses Proposition 3.9 to prove that the topological space underlying $\text{Ad}(X, \iota)$ coincides with the extended exploded tropicalization $\mathfrak{Top}(X, \iota)$:

Theorem 3.10. *Given any closed embedding $\iota : X \hookrightarrow Y_\Sigma$ of a K -variety X into a proper toric K -variety Y_Σ , there is a natural bijection*

$$\text{Ad}(X, \iota) \xrightarrow{\sim} \mathfrak{Trop}(X, \iota)$$

from the underlying set of the adic tropicalization to the exploded extended tropicalization.

3.5 Adic Tropicalization and Metrized Complexes

In this section, we describe the adic tropicalization of a generic hyperplane in the toric variety \mathbb{P}^2 in detail, and we make some remarks about the relationship between adic tropicalizations and the metrized complexes of Amini and Baker [1].

3.5.1 Tropical Complexes and Metrized Complexes

Fix a K -variety X and a closed embedding $\iota : X \hookrightarrow Y_\Sigma$ into a toric variety. Each complete fan Δ in $N_{\mathbb{R}} \times \mathbb{R}_{\geq 0}$, with recession fan $\text{rec}(\Delta) = \Sigma$, determines a formal R -scheme \mathfrak{X}_Δ . As explained in Sect. 2.2.2, the associated adic space $\mathfrak{X}_\Delta^{\text{ad}} \cong X^{\text{ad}}$ comes with a specialization morphism

$$\text{sp}_{\mathfrak{X}_\Delta} : X^{\text{ad}} \longrightarrow \mathfrak{X}_\Delta.$$

Because the Berkovich analytification X^{an} is the maximal Hausdorff quotient $h : X^{\text{ad}} \twoheadrightarrow X^{\text{an}}$, we also have a map

$$q : X^{\text{ad}} \xrightarrow{h} X^{\text{an}} \xrightarrow{\text{trop}} \text{Trop}(X, \iota).$$

Define the *tropical complex (of varieties) associated with Δ* , denoted $\mathfrak{C}\mathfrak{X}_\Delta$, to be the image, inside the product of topological spaces $\text{Trop}(X, \iota) \times \mathfrak{X}_\Delta$, of the product map

$$q \times \text{sp}_{\mathfrak{X}_\Delta} : X^{\text{ad}} \longrightarrow \text{Trop}(X, \iota) \times \mathfrak{X}_\Delta.$$

Example 3.11 (Metrized Complex Associated with a Generic Hyperplane in \mathbb{P}^2). For a concrete example, let $Y_\Sigma = \mathbb{P}^2$ with its standard toric structure, let $\iota : X \hookrightarrow Y$ be the closed embedding of the hyperplane $X = V(x + y + 1)$ into \mathbb{P}^2 , and let C be any polyhedral complex decomposing $N_{\mathbb{R}} = \mathbb{R}^2$ such that the intersection $C_X \stackrel{\text{def}}{=} C \cap \text{Trop}(X, \iota)$ is the one pictured at left in Fig. 6 below. Then our tropical complex $\mathfrak{C}\mathfrak{X}_{\Delta_C}$ is the *metrized complex (of curves)*, in the sense of Amini and Baker [1, Sect. 1.2], pictured at right in Fig. 6.




Fig. 6 A Γ -rational polyhedral decomposition of $\text{Trop}(X)$ induced by a convex polyhedral decomposition C of $N_{\mathbb{R}}$, and its associated tropical complex $\mathcal{C}\mathfrak{X}_{\Delta_C}$

3.5.3 The General 1-Dimensional Case and Metrized Complexes

Consider the case where our embedding $\iota : X \hookrightarrow Y_{\Sigma}$ is a closed embedding of a smooth curve X into the toric variety Y_{Σ} . Recall that ι is *schön* if the initial degeneration X_v is smooth at every vector v in the tropicalization of each torus orbit in Y_{Σ} .

We claim if Y_{Σ} is proper and ι is schön, then each tropical complex $\mathcal{C}\mathfrak{X}_{\Delta}$ is a metrized complex (of curves) in the sense of Amini and Baker [1, Sect. 1.2]. Indeed, Lemma 3.9, in conjunction with the schön condition, implies that if δ is a cone in Δ whose intersection $\delta \cap (\text{Trop}(X, \iota) \times \{1\})$ is a (necessarily Γ -rational) point $(v, 1) \in \text{Trop}(X, \iota) \times \{1\}$, then the k -scheme underlying \mathfrak{X}_{δ} is the nonsingular curve $\text{in}_v X$. Any point in X^{ad} that gets mapped to v under $q : X^{\text{ad}} \rightarrow \text{Trop}(X, \iota)$ goes to a point of this nonsingular curve $\text{in}_v X$ under the specialization map $\text{sp}_{\mathfrak{X}_{\Delta}} : X^{\text{ad}} \rightarrow \mathfrak{X}_{\Delta}$. Similarly, Lemma 3.9 and the schön condition imply that if $\delta \cap (\text{Trop}(X, \iota) \times \{1\})$ is an edge connecting points $(u, 1)$ and $(v, 1)$ in $\text{Trop}(X, \iota) \times \{1\}$, then the k -scheme underlying \mathfrak{X}_{δ} is a k -curve with two smooth components meeting at a point, one of the components containing $\text{in}_u X$ and the other containing $\text{in}_v X$. The points of X^{ad} mapping to $\mathfrak{X}_{\delta} - (\text{in}_u X \cup \text{in}_v X)$ under the specialization map $\text{sp}_{\mathfrak{X}_{\Delta}}$ are mapped to interior points of the edge corresponding to δ inside $\text{Trop}(X, \iota)$.

In this way, we can describe the tropical complex $\mathcal{C}\mathfrak{X}_{\Delta}$ using exactly the kind of gluing datum that Amini and Baker use to define a metrized complex in [1, Sect. 1.2].

Example 3.12 (An Example of a 2-Dimensional Tropical Complex). Let $X = Y_{\Sigma} = \mathbb{P}^2$, where $\Sigma \subset N_{\mathbb{R}} = \mathbb{R}^2$ is the standard fan describing \mathbb{P}^2 as a toric variety. Let C be the polyhedral decomposition of $N_{\mathbb{R}} = \mathbb{R}^2$ pictured in Fig. 7 below. It determines a fan Δ_C in $N_{\mathbb{R}} \times \mathbb{R}_{\geq 0}$ with recession fan $\text{rec}(\Delta_C) = \Sigma$. Using Lemma 3.9, we can build the tropical complex $\mathcal{C}\mathfrak{X}_{\Delta_C}$ from the datum of the k -varieties underlying the formal models \mathfrak{X}_{δ_P} associated with each polygon P in C .




Fig. 7 The tropical complex $\mathfrak{C}\mathfrak{X}_{\Delta_C}$ in Example 3.12

The polyhedral complex C consists of three vertices, edges d_i and e_i for $1 \leq i \leq 3$, and 2-dimensional polygons P_j for $0 \leq j \leq 3$. The tropical complex $\mathfrak{C}\mathfrak{X}_{\Delta_C}$ is glued from a copy of \mathbb{P}_k^1 at each vertex, a product $D_i \times d_i$ with a 1-dimensional k -variety D_i associated with each edge d_i , $1 \leq i \leq 3$, and similarly for edges e_i , and then the three 2-dimensional polygons P_j , $1 \leq j \leq 3$, understood as products $\text{Spec } k \times P_j$. The data telling us how to glue these topological spaces comes from the incidence relations between the irreducible components of the k -variety underlying the formal model \mathfrak{X}_{Δ_C} .

Lemma 3.13. *If $\iota : X \hookrightarrow Y_\Sigma$ is a closed embedding into a proper toric variety, then there is a canonical bijection $\lim_{\leftarrow \text{REC}_{\Sigma, \iota}} \mathfrak{C}\mathfrak{X}_\Delta \xrightarrow{\sim} \text{Ad}(X, \iota)$.*

Proof. Each projection $\text{Trop}(X, \iota) \times \mathfrak{X}_\Delta \rightarrow \mathfrak{X}_\Delta$ induces a map $\mathfrak{C}\mathfrak{X}_\Delta \rightarrow \mathfrak{X}_\Delta$, and these maps give rise to a map

$$\lim_{\leftarrow \text{REC}_{\Sigma, \iota}} \mathfrak{C}\mathfrak{X}_\Delta \rightarrow \text{Ad}(X, \iota). \tag{16}$$

Fix a compatible system of points $x = (x_\Delta)_{\Delta \in \text{Rec}_{\Sigma, \iota}}$, where x_Δ lies in the k -variety underlying \mathfrak{X}_Δ . This system describes a point x in $\text{Ad}(X, \iota)$. Because $\mathfrak{X}_{\delta_1} \cap \mathfrak{X}_{\delta_2} = \emptyset$ in \mathfrak{X}_Δ whenever $\delta_1 \cap \delta_2 = \emptyset$ in $N_{\mathbb{R}} \times \mathbb{R}_{\geq 0}$, we know that the compatible system x has an associated point $v_x \in \text{Trop}(X, \iota)$, determined by the condition that v_x lie in the polygon P_δ associated with a cone δ in Δ whenever x_Δ lies in \mathfrak{X}_δ . Here P_δ denotes the unique polygon in $N_{\mathbb{R}}$ satisfying $P_\delta \times \{1\} = \delta \cap (N_{\mathbb{R}} \times \{1\})$. The very definition of v_x implies that if x' is any point of \mathfrak{X}^{ad} mapping to x under $\text{sp}_{\mathfrak{X}_\Delta} : X^{\text{ad}} \rightarrow \mathfrak{X}_\Delta$, then x' maps to v_x under the map $q : X^{\text{ad}} \rightarrow \text{Trop}(X, \iota)$. This implies subjectivity of the comparison map (16).

To see that (16) is injective, note that the argument of the previous paragraph implies that points given by pairs (x, u) and (y, v) in $\mathfrak{C}\mathfrak{X}_\Delta$ must map to distinct

points of $\text{Ad}(X, \iota)$ if their second coordinates $u, v \in \text{Trop}(X, \iota)$ differ. On the other hand, if $u = v$, but the inverse systems $y = (y_\Delta)_{\Delta \in \text{Rec}_{\Sigma, \iota}}$ and $x = (x_\Delta)_{\Delta \in \text{Rec}_{\Sigma, \iota}}$ differ, then they map to distinct points of $\text{Ad}(X, \iota)$. \square

Example 3.14 (The Adic Tropicalization of a Generic Hyperplane in \mathbb{P}^2). As in Example 3.11 above, let $\iota : X \hookrightarrow Y$ be the closed embedding of the hyperplane $X = V(x + y + 1)$ into \mathbb{P}^2 . Lemma 3.13 says that we can visualize $\text{Ad}(X, \iota)$ by understanding it as the object that results, in the inverse limit, when we bubble of a copy of \mathbb{P}_k^1 at every Γ -rational point in $\text{Trop}(X, \iota)$. The resulting topological space appears in Fig. 8 below.

3.5.7 Tropical “Type 5” Points and Inverse Systems of R -Models.

We take a moment to highlight a phenomenon that occurs in adic tropicalizations, related to the presence of “type 5” points in the adic affine line $(\mathbb{A}^1)^{\text{ad}}$.

Consider a polygonal interval $P_1 = \bullet \text{---} \bullet$ inside $N_{\mathbb{R}} = \mathbb{R}$. The k -variety underlying the model \mathfrak{X}_{P_1} consists of two copies of \mathbb{A}_k^1 intersecting at a single k -point x_1 as pictured in Fig. 9 below. If we subdivide P_1 by adding a single Γ -rational vertex in the interior of P_1 , then the k -variety underlying our formal R -scheme picks up a \mathbb{P}_k^1 component where x_1 used to be.

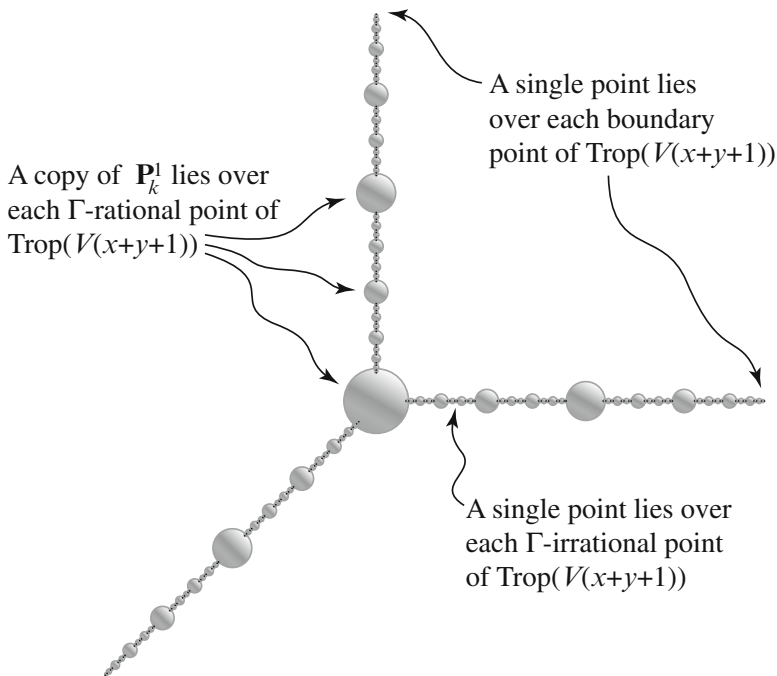


Fig. 8 Adic tropicalization of the hyperplane $V(x + y + 1)$ in \mathbb{P}^2 . Compare this to the exploded curve, in the sense of Parker [21], appearing in Fig. 1




Fig. 9 The formal R -schemes associated with several decompositions of the interval P_1 inside $N_{\mathbb{R}} = \mathbb{R}$. For each decomposition, the k -variety underlying the R -scheme \mathfrak{X}_v associated with the left endpoint v is a copy of $\mathbb{A}_k^1 - \{0\}$ that does not change from decomposition to decomposition. The nodal point x_i that partially compactifies $\mathbb{A}_k^1 - \{0\}$ in each of the R -schemes “stays fixed” as further decompositions move every \mathbb{P}_k^1 component further and further from \mathfrak{X}_v

Clearly \mathfrak{X}_v sits in the inverse limit of R -models associated with all subdivisions of the interval, but the inverse limit also contains a point in the closure of \mathfrak{X}_v , which does *not* map to the component of any model associated with a segment *not* containing v . The sequence of points x_i in Fig. 9 form an inverse system, and hence determine a single point x in the adic tropicalization $\text{Ad}(\mathbb{P}^1)$.

4 The Adic Limit Theorem

In this final Sect. 4, we explain that adic tropicalizations satisfy a form of “analytification is the limit of all tropicalizations,” namely, that one can recover the Huber analytification X^{ad} of any closed subvariety of a proper toric K -variety, along with the sheaf of power bound sections $\mathcal{O}_{X^{\text{ad}}}^{\circ}$ on X^{ad} , from an inverse limit of adic tropicalizations of X . The exact statement appears in Theorem 4.4 and Proposition 4.8 below. We begin with a brief review of the corresponding theorem in the setting of Berkovich analytic spaces, as proved in [23] and [6].

4.1 The Limit Theorem in the Berkovich Setting

Since each multiplicative seminorm $K[U_\sigma] \rightarrow \mathbb{R}_{\geq 0}$ induces a homomorphism $S_\sigma \rightarrow \mathbb{R} \sqcup \{\infty\}$ of semigroups, each closed embedding $\iota : X \hookrightarrow Y_\Sigma$ comes with a continuous map of topological spaces

$$\pi_\iota : X^{\text{an}} \rightarrow \text{Trop}(X, \iota),$$

called the *tropicalization map* associated with the closed embedding ι .

If $f_\phi : Y_\Sigma \rightarrow Y_{\Sigma'}$ is a toric morphism, necessarily induced by morphisms $\phi : \Sigma \rightarrow \Sigma'$ of fans, and if $\iota : X \hookrightarrow Y_\Sigma$ and $\iota' : X \hookrightarrow Y_{\Sigma'}$ are closed embeddings for which the diagram

$$\begin{array}{ccc}
 & & Y_\Sigma \\
 & \nearrow \iota & \downarrow f_\phi \\
 X & \hookrightarrow & \\
 & \searrow \iota' & \\
 & & Y_{\Sigma'}
 \end{array}
 \tag{17}$$

commutes, then the corresponding tropicalization maps π_ι and $\pi_{\iota'}$ fit into a commutative diagram

$$\begin{array}{ccc}
 & & \text{Trop}(X, \iota) \\
 & \nearrow \pi_\iota & \downarrow \text{Trop}(f_\phi) \\
 X^{\text{an}} & & \\
 & \searrow \pi_{\iota'} & \\
 & & \text{Trop}(X, \iota')
 \end{array}
 \tag{18}$$

Let \mathcal{S} denote any small category whose set of objects consists of some family of closed embeddings $\iota : X \hookrightarrow Y_\Sigma$, and in which each homset $\text{Hom}_{\mathcal{S}}(\iota, \iota')$ consists of all commutative diagrams of the above form (17). Then the system of tropicalization maps π_ι associated objects and morphisms in \mathcal{S} induce a map

$$\pi \stackrel{\text{def}}{=} \varprojlim_{\mathcal{S}} \pi_\iota : X^{\text{an}} \rightarrow \varprojlim_{\mathcal{S}} \text{Trop}(X, \iota) \tag{19}$$

We call this map the *Berkovich comparison map*.

In [23], S. Payne showed that if X admits at least one closed embedding into a quasi-projective toric variety Y_Σ , and if \mathcal{S} is the system of all closed embeddings into quasi-projective toric varieties, then the Berkovich comparison map (19) is a homeomorphism. In [6], Gross, Payne, and the present author extended this result to the case of any K -scheme admitting at least one closed embedding into a toric variety. This result is a consequence of the following Theorem 4.1, which describes

conditions under which an arbitrary system \mathcal{S} of closed toric embeddings gives rise to a homeomorphism (19):

Theorem 4.1 ([6, Theorem 1.1]). *If the system \mathcal{S} satisfies conditions (C1) and (C2) below, then the Berkovich comparison map (19) is a homeomorphism:*

- (C1) *If $\iota : X \hookrightarrow Y_\Sigma$ and $\iota' : X \hookrightarrow Y_{\Sigma'}$ are closed embeddings in \mathcal{S} , then their product $\iota \times \iota' : X \hookrightarrow Y_{\Sigma \times \Sigma'}$ is also in \mathcal{S} and*
- (C2) *There exists a finite affine open cover $\{ U_i \hookrightarrow X \}$ such that for each nonzero regular function $f \in K[U_i]$, there is some closed embedding $\iota : X \hookrightarrow Y_\Sigma$ in \mathcal{S} that realizes U_i as the preimage of a torus-invariant open affine, and realizes f as the pullback of a monomial.*

4.2 Adic Tropicalization Morphisms

Fix a K -variety X , and let $\iota : X \hookrightarrow Y_\Sigma$ be a closed embedding into a proper toric variety Y_Σ . As explained in Sect. 3.3, each Γ -admissible fan Δ in $N_{\mathbb{R}} \times \mathbb{R}_{\geq 0}$, satisfying

$$\text{supp}(\Delta) = N_{\mathbb{R}} \times \mathbb{R}_{\geq 0} \quad \text{and} \quad \Delta \cap (N_{\mathbb{R}} \times \{0\}) = \Sigma \times \{0\}, \tag{20}$$

gives rise to an admissible formal model $\mathfrak{X}_{(\iota, \Delta)}$ of X . By 2.2.2, this model $\mathfrak{X}_{(\iota, \Delta)}$ comes with a morphism of locally topologically ringed spaces

$$\text{sp}_{\mathfrak{X}_{(\iota, \Delta)}} : (X^{\text{ad}}, \mathcal{O}_{X^{\text{ad}}}^\circ) \longrightarrow (\mathfrak{X}_{(\iota, \Delta)}, \mathcal{O}_{\mathfrak{X}_{(\iota, \Delta)}}). \tag{21}$$

Our very construction of the adic tropicalization $\text{Ad}(X, \iota)$ as an inverse limit implies that as Δ ranges over all Γ -admissible fans in $N_{\mathbb{R}} \times \mathbb{R}_{\geq 0}$ satisfying the twin conditions (20), the resulting specialization morphisms (21) give rise to a single morphism of locally topologically ringed spaces

$$\varpi_\iota \stackrel{\text{def}}{=} \varprojlim_{\text{REC}_{\Sigma, \iota}} \text{sp}_{(\iota, \Delta)} : (X^{\text{ad}}, \mathcal{O}_{X^{\text{ad}}}^\circ) \longrightarrow (\text{Ad}(X, \iota), \mathcal{O}_{\text{Ad}(X, \iota)}). \tag{22}$$

This morphism (22) is the adic counterpart to the tropicalization map

$$\pi_\iota : X^{\text{an}} \longrightarrow \text{Trop}(X, \iota),$$

and so we refer to (22) as the *adic tropicalization morphism at ι* .

Definition 4.2. A *system of embeddings* for X is any category \mathcal{S} such that $\text{ob } \mathcal{S}$ is some class of closed embeddings $\iota : X \hookrightarrow Y_\Sigma$ into proper toric varieties,

and where the set of morphisms between two such embeddings is the set of all commutative diagrams

$$\begin{array}{ccc}
 & & Y_\Sigma \\
 X \hookrightarrow & \xrightarrow{i} & \\
 & \searrow i' & \downarrow f_\phi \\
 & & Y_{\Sigma'}
 \end{array} \tag{23}$$

where f_ϕ is any torus-equivariant morphism induced by a morphism $\phi : \Sigma \rightarrow \Sigma'$ of fans.

Remark 4.3. If Σ and Σ' are complete fans in $N_\mathbb{R}$ and $N'_\mathbb{R}$, then for each fan Δ' in $N'_\mathbb{R} \times \mathbb{R}_{\geq 0}$ satisfying the twin conditions (20), we can refine any fan Δ in $N_\mathbb{R} \times \mathbb{R}_{\geq 0}$ to a fan Δ_+ such that the product map

$$f_\phi \times \text{id} : N_\mathbb{R} \times \mathbb{R}_{\geq 0} \rightarrow N'_\mathbb{R} \times \mathbb{R}_{\geq 0}$$

induces a morphism of fans $f_\phi \times \text{id} : \Delta_+ \rightarrow \Delta'$. Thus if \mathcal{S} is a system of embeddings for X , then each morphism (23) induces a morphism of locally topologically ringed spaces

$$\text{Ad}(X, f_\phi) : \text{Ad}(X, i) \rightarrow \text{Ad}(X, i'). \tag{24}$$

These morphisms (24) give the assignment $(i : X \hookrightarrow Y_\Sigma) \mapsto \text{Ad}(X, i)$ the structure of a functor

$$\text{Ad}(X, -) : \mathcal{S} \rightarrow \mathbf{LTRS}.$$

The inverse limit of the adic tropicalization morphisms (22) over \mathcal{S} provides us with a morphism of locally topologically ringed spaces

$$\lim_{\leftarrow i \in \mathcal{S}} \varpi_i : (X^{\text{ad}}, \mathcal{O}_{X^{\text{ad}}}^\circ) \rightarrow \lim_{\leftarrow i \in \mathcal{S}} (\text{Ad}(X, i), \mathcal{O}_{\text{Ad}(X, i)}), \tag{25}$$

an adic counterpart to the comparison map (19) that appears in the context of Berkovich analytic spaces.

Theorem 4.4. *Let X be a proper variety over an algebraically closed, non-trivially valued non-Archimedean field K , and let \mathcal{S} be a system of embeddings for X . If \mathcal{S} satisfies conditions (C0), (C1), and (C2) below, then the adic comparison morphism (25) is an isomorphism of locally topologically ringed spaces.*

- (C0) *If $i : X \hookrightarrow Y_\Sigma$ is a closed embedding in \mathcal{S} , with \mathbb{T} the dense torus in Y_Σ , then any translation of i by a K -point of \mathbb{T} is again in \mathcal{S} ;*
- (C1) *If $i : X \hookrightarrow Y_\Sigma$ and $i' : X \hookrightarrow Y_{\Sigma'}$ are closed embeddings in \mathcal{S} , then their product $i \times i' : X \hookrightarrow Y_{\Sigma \times \Sigma'}$ is also in \mathcal{S} ; and*

(C2) *There exists a finite affine open cover $\{ U_i \hookrightarrow X \}$ such that for each nonzero regular function $f \in K[U_i]$, there is a closed embedding $\iota : X \hookrightarrow Y_\Sigma$ in \mathcal{S} that realizes U_i as the preimage of a torus-invariant open affine, and realizes f as the pullback of a monomial.*

Corollary 4.5 (Adic Limit Theorem on Point Sets). *If \mathcal{S} satisfies conditions (C0) through (C2), then the point set underlying the Huber analytification of X is in natural bijection with the inverse limit of all extended exploded tropicalizations of X :*

$$|X^{\text{ad}}| \cong \varprojlim_{\iota \in \mathcal{S}} \mathfrak{Top}(X, \iota)$$

Proof. This follows immediately from Theorems 3.10 and 4.4. □

Remark 4.6. One proves Theorem 4.4 by showing that the system of all Gubler models of X , for all closed embeddings $\iota : X \hookrightarrow Y_\Sigma$ in proper toric embeddings, is cofinal in the system of all admissible formal R -models of X . The key technical step in the proof is Proposition 4.7 below. As in the statement of Theorem 4.4, let X be a proper K -variety over an algebraically closed, non-trivially valued non-Archimedean field K , and let \mathcal{S} be a system of embeddings for X satisfying conditions (C0), (C1), and (C2) of Theorem 4.4.

Proposition 4.7. *Every admissible formal R -model \mathfrak{X} of X is dominated by a formal Gubler model. More precisely, if \mathfrak{X} is any admissible formal R -model of X , then there exist*

- (i) *a closed embedding $\iota : X \hookrightarrow Y_\Sigma$ into the proper toric variety associated with some complete fan Σ in some \mathbb{R} -vector space $N_{\mathbb{R}}$,*
- (ii) *a Γ -admissible fan Δ in $N_{\mathbb{R}} \times \mathbb{R}_{\geq 0}$ satisfying the twin conditions (20),*

such that the formal Gubler model $\mathfrak{X}_{(\iota, \Delta)}$ associated with the pair (ι, Δ) admits a proper morphism

$$\mathfrak{X}_{(\iota, \Delta)} \twoheadrightarrow \mathfrak{X}$$

of formal R -models of X .

To close, we give several examples of systems of embeddings that satisfy the hypotheses of Theorem 4.4, so that the comparison morphism (25) becomes an isomorphism.

Proposition 4.8. *Let X be a K -variety satisfying the two-point condition (A2) appearing in Sect. 3.4 above. Then in each of the following situations, the system \mathcal{S} satisfies conditions (C0), (C1), and (C2) of Theorem 4.4:*

- (i) *X proper, \mathcal{S} the category of all closed embeddings of X into proper toric varieties;*

- (ii) X projective, \mathcal{S} the category of all closed embeddings of X into projective toric varieties; and
- (iii) X smooth and proper, \mathcal{S} the category of all closed embeddings of X into smooth, proper toric varieties.

Proof. Case (ii) follows from [23, Lemma 4.3]. Cases (i) and (iii) follow from the construction described in the proof of [6, Theorem 4.2], using the fact that we can construct the toric embedding with the requisite property in each of these cases, as explained in [29, Theorem A]. \square

Acknowledgements The author would like to thank the Simons Foundation and the organizers of the Simons Symposium on non-Archimedean and Tropical Geometry for bringing together a fantastic group of researchers for an incredibly stimulating week of mathematics. Special thanks go to S. Payne, as this paper is a report on joint work with him. The author thanks M. Baker, F. Baldassarri, W. Gubler, and D. Ranganathan for helpful remarks. Finally, the author thanks the reviewer for comments that greatly improved this paper.

The author is supported by NSF RTG grant DMS-0943832.

References

1. Amini, O., Baker, M.: Linear series on metrized complexes of algebraic curves. arXiv: 1204.3508 (2012)
2. Berkovich, V.G.: Spectral Theory and Analytic Geometry Over Non-Archimedean Fields. Mathematical Surveys and Monographs, vol. 33. American Mathematical Society, Providence, RI (1990)
3. Bosch, S.: Lectures on Formal and Rigid Geometry. Lecture Notes in Mathematics, vol. 2105. Springer, Cham (2014)
4. Cox, D.A., Little, J.B., Schenck, H.K.: Toric Varieties. Graduate Studies in Mathematics, vol. 124. American Mathematical Society, Providence, RI (2011)
5. Foster, T., Payne, S.: Limits of tropicalizations II: structure sheaves, adic spaces and cofinality of Gubler models (in preparation)
6. Foster, T., Gross, P., Payne, S.: Limits of tropicalizations. *Isr. J. Math.* **201**(2), 835–846 (2014)
7. Gillam, W.D.: Localization of ringed spaces. *Adv. Pure Math.* **1**(5), 250–263 (2011)
8. Gross, M.: Tropical Geometry and Mirror Symmetry. CBMS Regional Conference Series in Mathematics, vol. 114. American Mathematical Society, Providence, RI (2011); Published for the Conference Board of the Mathematical Sciences, Washington, DC. <http://www.ams.org/mathscinet/search/publdoc.html> (2011)
9. Gross, M., Siebert, B.: Affine manifolds, log structures, and mirror symmetry. *Turk. J. Math.* **27**(1), 33–60 (2003)
10. Gubler, W.: A guide to tropicalizations. In: Algebraic and Combinatorial Aspects of Tropical Geometry. Contemporary Mathematics, vol. 589, pp. 125–189. American Mathematical Society, Providence, RI (2013)
11. Gubler, W., Soto, A.: Classification of normal toric varieties over a valuation ring of rank one. arXiv: 1303.1987 (2013)
12. Hartshorne, R.: Algebraic geometry. Graduate Texts in Mathematics, No. 52. Springer, New York (1977)
13. Huber, R.: A generalization of formal schemes and rigid analytic varieties. *Math. Z.* **217**(4), 513–551 (1994)

14. Huber, R.: *Étale Cohomology of Rigid Analytic Varieties and Adic Spaces*. Aspects of Mathematics, vol. E30. Friedr. Vieweg & Sohn, Braunschweig (1996)
15. Kempf, G., Knudsen, F.F., Mumford, D., Saint-Donat, B.: *Toroidal Embeddings. I*. Lecture Notes in Mathematics, vol. 339. Springer, Berlin (1973)
16. Kontsevich, M., Soibelman, Y.: Homological mirror symmetry and torus fibrations. In: *Symplectic Geometry and Mirror Symmetry* (Seoul, 2000), pp. 203–263. World Science Publisher, River Edge, NJ (2001)
17. Kontsevich, M., Soibelman, Y.: Affine structures and non-Archimedean analytic spaces. In: *The Unity of Mathematics*. Progress in Mathematics, vol. 244, pp. 321–385. Birkhäuser, Boston, MA (2006)
18. Maclagan, D., Sturmfels, B.: *Introduction to Tropical Geometry*. Mathematical Surveys and Monographs, vol. 161. American Mathematical Society, Providence, RI (2015)
19. Parker, B.: Exploded fibrations. In: *Proceedings of Gökova Geometry-Topology Conference 2006*, Gökova Geometry/Topology Conference (GGT), Gökova, pp. 52–90 (2007)
20. Parker, B.: Gromov Witten invariants of exploded manifolds. arXiv: 1102.0158 (2011)
21. Parker, B.: Exploded manifolds. *Adv. Math.* **229**(6), 3256–3319 (2012)
22. Parker, B.: Log geometry and exploded manifolds. *Abh. Math. Semin. Univ. Hambg.* **82**(1), 43–81 (2012)
23. Payne, S.: Analytification is the limit of all tropicalizations. *Math. Res. Lett.* **16**(3), 543–556 (2009)
24. Payne, S.: Fibers of tropicalization. *Math. Z.* **262**(2), 301–311 (2009)
25. Payne, S.: Erratum to: Fibers of tropicalization. *Math. Z.* **272**(3–4), 1403–1406 (2012)
26. Scholze, P.: Perfectoid spaces. *Publ. Math. Inst. Hautes Études Sci.* **116**, 245–313 (2012)
27. van der Put, M., Schneider, P.: Points and topologies in rigid geometry. *Math. Ann.* **302**(1), 81–103 (1995)
28. Wedhorn, T.: Adic spaces. <http://www2.math.uni-paderborn.de/fileadmin/Mathematik/People/wedhorn/Lehre/AdicSpaces.pdf> (2012)
29. Włodarczyk, J.: Embeddings in toric varieties and prevarieties. *J. Algebraic Geom.* **2**(4), 705–726 (1993)

Degeneration of Linear Series from the Tropical Point of View and Applications

Matthew Baker and David Jensen

Abstract We discuss linear series on tropical curves and their relation to classical algebraic geometry, describe the main techniques of the subject, and survey some of the recent major developments in the field, with an emphasis on applications to problems in Brill–Noether theory and arithmetic geometry.

Keywords Chip firing • Tropical curves • Brill–Noether theory

1 Introduction

Algebraic curves play a central role in the field of algebraic geometry. Over the past century, curves have been the focus of a significant amount of research, and despite being some of the most well-understood algebraic varieties, there are still many important open questions. The goal of classical Brill–Noether theory is to study the geometry of a curve C by examining all of its maps to projective space, or equivalently the existence and behavior of all line bundles on C . Thus, we have classical results such as Max Noether’s Theorem [12, p. 117] and the Enriques–Babbage Theorem [13] that relate the presence of linear series on a curve to its geometric properties. A major change in perspective occurred during the twentieth century, as the field shifted from studying *fixed* to *general* curves—that is, general points in the moduli space of curves M_g . Many of the major results in the field, such as the Brill–Noether [66] and Gieseker–Petri [65] theorems, remained open for nearly a century as they awaited this new point of view.

A major milestone in the geometry of general curves was the development of limit linear series by Eisenbud and Harris [58]. This theory allows one to study linear series on general curves by studying one-parameter degenerations where the central fiber is a special kind of singular curve, known as a curve of compact type.

M. Baker (✉)

School of Mathematics, Georgia Institute of Technology, Atlanta, GA 30332-0160, USA
e-mail: mbaker@math.gatech.edu

D. Jensen

Department of Mathematics, University of Kentucky, Lexington, KY 40506-0027, USA
e-mail: dave.jensen@uky.edu

One property of curves of compact type is that if they have positive genus then they must have components of positive genus. Shortly after the development of limit linear series, many researchers became interested in a different type of nodal curve, which has only rational components, and where the interesting geometric data is encoded by the combinatorics of how the components meet each other. Early examples using the so-called graph curves to establish properties of general curves include Bayer and Eisenbud's work on Green's conjecture for the general curve [28], and the Ciliberto–Harris–Miranda result on the surjectivity of the Wahl map [47].

Much like the theory of limit linear series does for curves of compact type, the recent development of tropical Brill–Noether theory provides a systematic approach to this kind of degeneration argument [5, 17, 19]. A major goal of this survey is to introduce the basic definitions and techniques of this theory, as well as describing some recent applications to the geometry of general curves and the behavior of Weierstrass points in degenerating families. Degeneration arguments also play a major role in arithmetic geometry, and we also survey how linear series on tropical curves can be used to study rational points on curves.

Here are just a few of the interesting theorems which have been proved in recent years with the aid of the theory of linear series on tropical curves.

1. *The Maximal Rank Conjecture for quadrics.* In [79], Jensen and Payne prove that for fixed g, r , and d , if C is a general curve of genus g and $V \subset \mathcal{L}(D)$ is a general linear series on C of rank r and degree d , then the multiplication map $\mu_2 : \text{Sym}^2 V \rightarrow \mathcal{L}(2D)$ is either injective or surjective.
2. *Uniform boundedness for rational points of curves of small Mordell–Weil rank.* In [81], Katz, Rabinoff, and Zureick–Brown prove that if C/\mathbb{Q} is a curve of genus g with Mordell–Weil rank at most $g - 3$, then $\#C(\mathbb{Q}) \leq 76g^2 - 82g + 22$. This is the first such bound depending only on the genus of C .
3. *Non-Archimedean equidistribution of Weierstrass points.* In [4], Amini proves that if C is an algebraic curve over $\mathbb{C}((t))$ and L is an ample line bundle on C , then the Weierstrass points of $L^{\otimes n}$ become equidistributed with respect to the Zhang measure on the Berkovich analytic space C^{an} as n goes to infinity. This gives precise asymptotic information on the limiting behavior of Weierstrass points in degenerating one-parameter families.
4. *Mnev universality for the lifting problem for divisors on graphs.* In [40], Cartwright shows that if X is a scheme of finite type over $\text{Spec } \mathbb{Z}$, there exist a graph G and a rank 2 divisor D_0 on G such that, for any infinite field k , there are a curve C over $k((t))$ and a rank 2 divisor D on C tropicalizing to G and D_0 , respectively, if and only if X has a k -point.

We will discuss the proofs of these and other results after going through the foundations of the basic theory. To accommodate readers with various interests, this survey is divided into three parts. The first part covers the basics of tropical Brill–Noether theory, with an emphasis on combinatorial aspects and the relation to classical algebraic geometry. The second part covers more advanced topics, including the non-Archimedean Berkovich space perspective, tropical moduli spaces, and the theory of metrized complexes. Each of these topics is an important part of the

theory, but is not strictly necessary for many of the applications discussed in Part 3, and the casual reader may wish to skip Part 2 on the first pass. The final part covers applications of tropical Brill–Noether theory to problems in algebraic and arithmetic geometry. For the most part, the sections in Part 3 are largely independent of each other and can be read in any order. Aside from a few technicalities, the reader can expect to follow the applications in Sects. 9, 10, and 11, as well as most of Sect. 12, without reading Part 2.

Part 1: Introductory Topics

2 Jacobians of Finite Graphs

2.1 Degeneration of Line Bundles in One-Parameter Families

A recurring theme in the theory of linear series on curves is that it is very important to understand the behavior of line bundles on generically smooth one-parameter families of curves. One of the key facts about such families is the semistable reduction theorem, which asserts that after a finite base change, one can guarantee that the singularities of the family are “as nice as possible,” i.e., the total space is regular and the fibers are reduced and have only nodal singularities (see [72, Chap. 3C] or [94, Sect. 10.4]). The questions we want to answer about such families are all local on the base, so it is convenient to consider the following setup. Let R be a discrete valuation ring with field of fractions K and algebraically closed residue field κ , let C be a smooth proper and geometrically integral curve over K , and let \mathcal{C} be a *regular strongly semistable model* for C , that is, a proper flat R -scheme with general fiber C satisfying:

1. The total space \mathcal{C} is regular.
2. The central fiber C_0 of \mathcal{C} is *strongly semistable*, i.e., the irreducible components of C_0 are all smooth and C_0 has only nodes as singularities.¹

By the semistable reduction theorem, a regular strongly semistable model for C always exists after passing to a finite extension of K .

Let L be a line bundle on the general fiber C . Because the total space \mathcal{C} is regular, there exists an extension \mathcal{L} of L to the family \mathcal{C} . One can easily see, however, that this extension is not unique—one can obtain other extensions by twisting the components of the central fiber. More concretely, if \mathcal{L} is an extension of L and $Y \subset C_0$ is an irreducible component of the central fiber, then $\mathcal{L}(Y) = \mathcal{L} \otimes \mathcal{O}_{\mathcal{C}}(Y)$ is also an extension of L .

¹We say C_0 is *semistable* if it is reduced and has only nodes as singularities. The term *strongly semistable* is not completely standard.

To understand the effect of the twisting operation, we consider the *dual graph* of the central fiber C_0 . One constructs this graph by first assigning a vertex v_Z to each irreducible component Z of C_0 , and then drawing an edge between two vertices for every node at which the corresponding components intersect.

Example 2.1. If C_0 is a union of m general lines in \mathbb{P}^2 , its dual graph is the complete graph on m vertices. Indeed, every pair of lines meets in one point, so between every pair of vertices in the dual graph, there must be an edge.

Example 2.2. If $C_0 \subset \mathbb{P}^1 \times \mathbb{P}^1$ is a union of a lines in one ruling and b lines in the other ruling, then the dual graph is the complete bipartite graph $K_{a,b}$. This is because a pair of lines in the same ruling do not intersect, whereas a pair of lines in opposite rulings intersect in one point.

Example 2.3. Let C_0 be the union of the -1 curves on a del Pezzo surface of degree 5. In this case, the dual graph of C_0 is the well-known *Petersen graph*.

2.2 Divisors and Linear Equivalence on Graphs

In this paper, by a *graph* we will always mean a finite connected graph which is allowed to have multiple edges between pairs of vertices but is not allowed to have any loop edges. Given a graph G , we write $\text{Div}(G)$ for the free abelian group on the vertices of G . An element D of $\text{Div}(G)$ is called a *divisor* on G , and is written as a formal sum

$$D = \sum_{v \in V(G)} D(v)v,$$

where the coefficients $D(v)$ are integers. The *degree* of a divisor D is defined to be the sum

$$\text{deg}(D) = \sum_{v \in V(G)} D(v).$$

Returning now to our family of curves, let us fix for a moment an extension \mathcal{L} of our line bundle L . We define a corresponding divisor $\text{mdeg}(\mathcal{L})$ on G , called the *multidegree* of \mathcal{L} , by the formula

$$\text{mdeg}(\mathcal{L}) = \sum_Z (\text{deg } \mathcal{L}|_Z) v_Z,$$

where the sum is over all irreducible components Z of C_0 . The quantity $\text{deg } \mathcal{L}|_Z$ can also be interpreted as the intersection multiplicity of \mathcal{L} with Z considered as a (vertical) divisor on the surface \mathcal{C} . Note that the degree of $\text{mdeg}(\mathcal{L})$ is equal to $\text{deg}(L)$.

We now ask ourselves how $\text{mdeg}(\mathcal{L})$ changes if we replace \mathcal{L} by a different extension. Since any two such extensions differ by a sequence of twists by components of the central fiber, it suffices to study the effect of twisting by one such component Y . Given a vertex v of a graph, let $\text{val}(v)$ denote its valence. As the central fiber C_0 is a principal divisor on \mathcal{C} , we have

$$\text{mdeg}(\mathcal{L}(Y)) = \text{mdeg}(\mathcal{L}) + \sum_Z (Y \cdot Z) v_Z$$

with

$$Y \cdot Z = \begin{cases} -\text{val}(v_Y) & \text{if } Z = Y \\ |Z \cap Y| & \text{if } Z \neq Y \end{cases}.$$

The corresponding operation on the dual graph is known as a *chip-firing move*. This is because we may think of a divisor on the graph as a configuration of chips (and anti-chips) on the vertices. In this language, the effect of a chip-firing move is that a vertex v “fires” one chip along each of the edges emanating from v . This decreases the number of chips at v by the valence of v , and increases the number of chips at each of the neighbors w of v by the number of edges between v and w .

Motivated by these observations, we say that two divisors D and D' on a graph G are *equivalent*, and we write $D \sim D'$, if one can be obtained from the other by a sequence of chip-firing moves. We define the *Picard group* $\text{Pic}(G)$ of G to be the group of divisors on G modulo equivalence. Note that the degree of a divisor is invariant under chip-firing moves, so there is a well-defined homomorphism

$$\text{deg} : \text{Pic}(G) \rightarrow \mathbb{Z}.$$

We will refer to the kernel $\text{Pic}^0(G)$ of this map as the *Jacobian* $\text{Jac}(G)$ of the graph G . This finite abelian group goes by many different names in the mathematical literature—in combinatorics, it is commonly referred to as the *sandpile group* or the *critical group* of G (see, for example, [30, 56, 92, 109]).

Remark 2.4. There is a tremendous amount of combinatorial literature concerning Jacobians of graphs. As our focus is on applications in algebraic geometry, however, we will not go into many details here—we refer the reader to [19, 76] and the references therein. We cannot resist mentioning one remarkable fact, however: the cardinality of $\text{Jac}(G)$ is the number of *spanning trees* in G . (This is actually a disguised form of Kirchhoff’s celebrated *Matrix-Tree Theorem*.)

Our discussion shows that for any two extensions of the line bundle L , the corresponding multidegrees are equivalent divisors on the dual graph G . There is therefore a well-defined degree-preserving map

$$\text{Trop} : \text{Pic}(C) \rightarrow \text{Pic}(G),$$

which we refer to as the *specialization* or *tropicalization* map.

For later reference, we note that there is a natural homomorphism $\rho : \text{Div}(C) \rightarrow \text{Div}(G)$ defined by setting $\rho(D) = \text{mdeg}(D)$, where \mathcal{D} is the Zariski closure of D in \mathcal{C} . The map ρ takes principal (resp., effective) divisors to principal (resp., effective) divisors, and the map $\text{Jac}(C) \rightarrow \text{Jac}(G)$ induced by ρ coincides with Trop (cf. [17, Sect. 2.1]).

It is useful to reformulate the definition of equivalence of divisors on G as follows. For any function $f : V(G) \rightarrow \mathbb{Z}$, we define

$$\text{ord}_v(f) := \sum_{w \text{ adjacent to } v} f(v) - f(w)$$

and

$$\text{div}(f) := \sum_{v \in V(G)} \text{ord}_v(f)v.$$

Divisors of the form $\text{div}(f)$ are known as *principal* divisors, and two divisors are equivalent if and only if their difference is principal. The reason is that the divisor $\text{div}(f)$ can be obtained, starting with the 0 divisor, by firing each vertex v exactly $f(v)$ times.

2.3 Limit Linear Series and Néron Models

The Eisenbud–Harris theory of limit linear series focuses on the case where the dual graph of C_0 is a tree. In this case, the curve C_0 is said to be of *compact type*.² Although they would not have stated it this way, a key insight of the Eisenbud–Harris theory is that the Jacobian of a tree is trivial. Given a line bundle \mathcal{L} of degree d on \mathcal{C} , if the dual graph of C_0 is a tree, then one can repeatedly twist to obtain a line bundle with any degree distribution summing to d on the components of the central fiber. In particular, given a component $Y \subset C_0$, there exists a twist \mathcal{L}_Y such that

$$\text{deg } \mathcal{L}_Y|_Z = \begin{cases} d & \text{if } Z = Y \\ 0 & \text{if } Z \neq Y \end{cases}.$$

This observation is the jumping-off point for the basic theory of limit linear series.

At the other end of the spectrum is the *maximally degenerate* case where all of the components of C_0 have genus 0. In this case, the first Betti number of the dual graph G (which we refer to from now on as the *genus* of G) is equal to the

²The reason for the name *compact type* is that the Jacobian of a nodal curve C_0 is an extension of the Jacobian of the normalization of C_0 by a torus of dimension equal to the first Betti number of the dual graph. It follows that the Jacobian of C_0 is compact if and only if its dual graph is a tree.

geometric genus of the curve and essentially all of the interesting information about degenerations of line bundles is combinatorial. At its core, tropical Brill–Noether theory studies the behavior of line bundles on the curve C using the combinatorics of their specializations to the graph G .

This discussion can also be understood in the context of Néron models. An important theorem of Raynaud [111] asserts (in the language of this paper) that the group Φ of connected components of the special fiber $\tilde{\mathcal{J}}$ of the Néron model of $J = \text{Jac}(C)$ is canonically isomorphic to $\text{Jac}(G)$, where G is the dual graph of the special fiber of \mathcal{C} . This result can be summarized by the commutativity of the following diagram:

$$\begin{array}{ccccccc}
 0 & \longrightarrow & \text{Prin}(C) & \longrightarrow & \text{Div}^0(C) & \longrightarrow & J(K) \longrightarrow 0 \\
 & & \downarrow \rho & & \downarrow \rho & & \downarrow \\
 0 & \longrightarrow & \text{Prin}(G) & \longrightarrow & \text{Div}^0(G) & \longrightarrow & \text{Jac}(G) = \Phi \longrightarrow 0
 \end{array} \tag{1}$$

where the right vertical arrow is the canonical quotient map

$$J(K) \rightarrow J(K)/\mathcal{J}^0(R) = \Phi.$$

The tropical approach to Brill–Noether theory and the approach via the theory of limit linear series are in some sense orthogonal. The former utilizes the component group Φ , or its analytic counterpart, the tropical Jacobian, whereas classical limit linear series are defined only when Φ is trivial. On the other hand, limit linear series involve computations in the compact part of $\tilde{\mathcal{J}}$, which is trivial in the maximally degenerate case.

Recent developments have led to a sort of hybrid of tropical and limit linear series that can be used to study degenerations of line bundles to arbitrary nodal curves. We will discuss these ideas in Sect. 8.

3 Jacobians of Metric Graphs

3.1 Behavior of Dual Graphs Under Base Change

The field of fractions of a DVR is never algebraically closed. For many applications, we will be interested in $\text{Pic}(C_{\bar{K}})$ rather than $\text{Pic}(C_K)$, and we must therefore study the behavior of the specialization map under base change.

Let K' be a finite extension of K , let R' be the valuation ring of K' , and let $C_{K'} = C \times_K K'$. An important issue is that the new total space $\mathcal{C}_{K'} = \mathcal{C} \times_K K'$ may not be regular; it can pick up singularities at the nodes of the central fiber. More specifically, if a point z on \mathcal{C} corresponding to a node of the central fiber has a local analytic equation of the form $xy = \pi$, where π is a uniformizer for

R , then a local analytic equation for z over R' will be $xy = (\pi')^e$, where π' is a uniformizer for R' and e is the ramification index of the extension K'/K . A standard computation shows that we can resolve such a singularity by a chain of $e - 1$ blowups. Repeating this procedure for each singular point of the special fiber, we obtain a regular strongly semistable model \mathcal{C}' for $C_{K'}$. The dual graph G' of the central fiber of \mathcal{C}' is obtained by subdividing each edge of the original dual graph G $e - 1$ times. In other words, if we assign a length of 1 to each edge of G , and a length of $\frac{1}{e}$ to each edge of G' , then G and G' are isomorphic as *metric graphs*.

A *metric graph* Γ is, roughly speaking, a finite graph G in which each edge e has been identified with a real interval I_e of some specified length $\ell_e > 0$. The points of Γ are the vertices of G together with all points in the relative interiors of the intervals I_e . More precisely, a metric graph is an equivalence class of finite edge-weighted graphs, where two weighted graphs G and G' give rise to the same metric graphs G and G' if they have a common length-preserving refinement. A finite weighted graph G representing the equivalence class of Γ is called a *model* for Γ .

Example 3.1. Let $K = \mathbb{C}((t))$, and consider the family $\mathcal{C} : xy = tz^2$ of smooth conics degenerating to a singular conic in \mathbb{P}^2 . The dual graph G of the central fiber consists of two vertices v, v' connected by a single edge. Let $D = P + Q$ be the divisor on the general fiber cut out by the line $y = x$. In homogeneous coordinates, we have $P = (\sqrt{t} : \sqrt{t} : 1)$ and $Q = (-\sqrt{t} : -\sqrt{t} : 1)$. Although the divisor D itself is K -rational, P and Q are not, and both points specialize to the node of the special fiber. It is not hard to check that $\rho(D) = v + v' \in \text{Div}(G)$. If $K' = \mathbb{C}((\sqrt{t}))$, then the total space of the family $\mathcal{C} \times_K K'$ has a singularity at the node of the central fiber. This singularity can be resolved by blowing up the node, and the dual graph G' of the new central fiber is a chain of 3 vertices connected by two edges, with the new vertex v'' corresponding to the exceptional divisor of the blowup (see Fig. 1). The base change D' of D to K' specializes to a sum of two smooth points on the exceptional divisor, and in particular $\rho_{\mathcal{C}'}(D') = 2v''$ is not the image of $\rho_{\mathcal{C}}(D)$ with respect to the natural inclusion map $\text{Div}(G) \hookrightarrow \text{Div}(G')$. Note, however, that $\rho_{\mathcal{C}'}(D')$ and $\rho_{\mathcal{C}}(D)$ are linearly equivalent on G' (this turns out to be a general phenomenon).

Remark 3.2. A similar example, with a semistable family of curves of genus 3, is given in [17, Sect. 4.4].




Fig. 1 The dual graph of the central fiber in Example 3.1 initially (*on the left*), and after base change followed by resolution of the singularity (*on the right*). If we give the segment *on the left* a length of 1 and each of the segments *on the right* a length of 1/2, then both weighted graphs are models for the same underlying metric graph Γ , which is a closed segment of length 1

There are two possible ways to address the lack of functoriality with respect to base change illustrated in Example 3.1. One is to only consider the induced maps on Picard groups, rather than divisors. The other is to replace $\rho : \text{Div}(C) \rightarrow \text{Div}(G)$ with a map $\text{Div}(C_{\bar{K}}) \rightarrow \text{Div}(\Gamma)$, where Γ is the metric graph underlying G , defined by first base-changing to an extension K' over which all points in the support of D are rational and then applying ρ . The most natural way to handle these simultaneous base changes and prove theorems about the resulting map is to work on the Berkovich analytification of C ; see Sect. 6 for details.

3.2 Divisors and Linear Equivalence on Metric Graphs

Let Γ be a metric graph. A *divisor* D on Γ is a formal linear combination

$$D = \sum_{v \in \Gamma} D(v)v$$

with $D(v) \in \mathbb{Z}$ for all $v \in \Gamma$ and $D(v) = 0$ for all but finitely many $v \in \Gamma$. Let $\text{PL}(\Gamma)$ denote the set of continuous, piecewise linear functions $f : \Gamma \rightarrow \mathbb{R}$ with integer slopes. The *order* $\text{ord}_v(f)$ of a function f at a point $v \in \Gamma$ is the sum of its incoming³ slopes along the edges containing v . As in the case of finite graphs, we write

$$\text{div}(f) := \sum_{v \in \Gamma} \text{ord}_v(f)v.$$

A divisor is said to be *principal* if it is of the form $\text{div}(f)$ for some $f \in \text{PL}(\Gamma)$, and two divisors D, D' are *equivalent* if $D - D'$ is principal. We let $\text{Prin}(\Gamma)$ denote the subgroup of $\text{Div}^0(\Gamma)$ (the group of degree-zero divisors on Γ) consisting of principal divisors. By analogy with the case of finite graphs, the group

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

is called the (tropical) *Jacobian* of Γ .

Example 3.3. The Jacobian of a metric tree Γ is trivial. To see this, note that given any two points $P, Q \in \Gamma$, there is a unique path from P to Q . One can construct a continuous, piecewise linear function f that has slope 1 along this path, and has slope 0 everywhere else. We then see that $\text{div}(f) = Q - P$, so any two points on the tree are equivalent (Fig. 2).

Example 3.4. A circle Γ is a torsor for its own Jacobian. To see this, fix a point $O \in \Gamma$. Given two points $P, Q \in \Gamma$, there exists a continuous, piecewise linear function f that has slope 1 on the interval from O and P , slope -1 on the interval from some 4th

³As a caution, some authors use the opposite sign convention.

Fig. 2 Slopes of a function f on a tree with $\text{div}(f) = Q - P$





Fig. 3 A function f on the circle with $\text{div}(f) = O + Q - (P + R)$



point R to Q , and slope 0 everywhere else. We then have $\text{div}(f) = O + Q - (P + R)$, so $P - Q \sim R - O$. It follows that every divisor of degree zero is equivalent to a divisor of the form $R - O$ for some point R . It is not difficult to see that the point R is in fact unique (Fig. 3).

Note that if two divisors are equivalent in the finite graph G , then they are also equivalent in the corresponding metric graph Γ , called the *regular realization* of G , in which every edge of G is assigned a length of 1. It follows that there is a natural inclusion $\iota : \text{Pic}(G) \hookrightarrow \text{Pic}(\Gamma)$. As we saw above, the multidegree of a line bundle L on $C_{\bar{K}}$ can be identified with a divisor on some subdivision of G , which can in turn be identified with a divisor on Γ . One can show that this yields a well-defined map

$$\text{Trop} : \text{Pic}(C_{\bar{K}}) \rightarrow \text{Pic}(\Gamma)$$

whose restriction to $\text{Pic}(C)$ coincides with the previously defined map

$$\text{Trop} : \text{Pic}(C) \rightarrow \text{Pic}(G)$$

via the inclusion ι .

One can see from this construction that principal divisors on C specialize to principal divisors on Γ . More precisely, there is a natural way to define a map

$$\text{trop} : \bar{K}(C)^* \rightarrow \text{PL}(\Gamma)$$

on rational functions such that

$$\text{Trop}(\text{div}(f)) = \text{div}(\text{trop}(f))$$

for every $f \in \bar{K}(C)^*$. This is known as the *Slope Formula*, cf. Theorem 6.4 below. We refer the reader to Sect. 6.3 for a formal definition of the map trop .

As already discussed in the remarks following Example 3.1, there is also a natural way to define a map $\text{Div}(C_{\overline{K}}) \rightarrow \text{Div}(\Gamma)$ which induces the map $\text{Trop} : \text{Pic}(C_{\overline{K}}) \rightarrow \text{Pic}(\Gamma)$. As with the map trop on rational functions, the map Trop on divisors is most conveniently described using Berkovich’s theory of non-Archimedean analytic spaces, and we defer a detailed discussion to Sect. 6.

3.3 The Tropical Abel–Jacobi Map

A 1-form on a graph G is an element of the real vector space generated by the formal symbols de , as e ranges over the oriented edges of G , subject to the relations that if e, e' represent the same edge with opposite orientations then $de' = -de$. After fixing an orientation on G , a 1-form $\omega = \sum \omega_e de$ is called *harmonic* if, for all vertices v , the sum $\sum \omega_e$ over the outgoing edges at v is equal to 0. Denote by $\Omega(G)$ the space of harmonic 1-forms on G . It is well-known that $\Omega(G)$ is a real vector space of dimension equal to the *genus* g of G , which can also be defined combinatorially as the number of edges of G minus the number of vertices of G plus one, or topologically as the dimension of $H_1(G, \mathbb{R})$.

If G and G' are models for the same metric graph Γ , then $\Omega(G')$ is canonically isomorphic to $\Omega(G)$. We may therefore define the space $\Omega(\Gamma)$ of harmonic 1-forms on Γ as $\Omega(G)$ for any weighted graph model G . (We also define the *genus* of a metric graph Γ to be the genus of any model for Γ .) Given an isometric path $\gamma : [a, b] \rightarrow \Gamma$, any harmonic 1-form ω on Γ pulls back to a classical 1-form on the interval, and we can thus define the integral $\int_{\gamma} \omega$. Note that the definition of a harmonic 1-form does not depend on the metric, but the integral $\int_{\gamma} \omega$ does.

Fix a base point $v_0 \in \Gamma$. For any point $v \in \Gamma$, the integral

$$\int_{v_0}^v \omega$$

is well-defined up to a choice of path from v_0 to v . This gives a map


$$AJ_{v_0} : \Gamma \rightarrow \Omega(\Gamma)^*/H_1(\Gamma, \mathbb{Z})$$

known as the *tropical Abel–Jacobi map*. Extending linearly to $\text{Div}(\Gamma)$ and then restricting to $\text{Div}^0(\Gamma)$, we obtain a map

$$AJ : \text{Div}^0(\Gamma) \rightarrow \Omega(\Gamma)^*/H_1(\Gamma, \mathbb{Z})$$

which does not depend on the choice of a base point. As in the classical case of Riemann surfaces, we have (cf. [101]):

Fig. 4 Two metric graphs Γ and Γ' of genus 2



Tropical Abel–Jacobi Theorem. *The map AJ is surjective and its kernel is precisely $\text{Prin}(\Gamma)$. Thus there is a canonical isomorphism*

$$AJ : \text{Div}^0(\Gamma)/\text{Prin}(\Gamma) \cong \Omega(\Gamma)^*/H_1(\Gamma, \mathbb{Z})$$

between the Jacobian of Γ and a g -dimensional real torus.

Example 3.5. We consider the two metric graphs of genus 2 pictured in Fig. 4. In the first case, we can choose a basis ω_1, ω_2 of harmonic 1-forms by assigning the integer 1 to one of the (oriented) loops, and the integer 0 to the other. We let η_1, η_2 be the elements of the dual basis. We then see from the tropical Abel–Jacobi theorem that

$$\text{Jac}(\Gamma) \cong \mathbb{R}^2/(\mathbb{Z}\eta_1 + \mathbb{Z}\eta_2).$$

A similar argument in the second case yields

$$\text{Jac}(\Gamma') \cong \mathbb{R}^2/ \left(\mathbb{Z}[\eta_1 + \frac{1}{2}\eta_2] + \mathbb{Z}[\frac{1}{2}\eta_1 + \eta_2] \right).$$

Although the Jacobians of the metric graphs Γ and Γ' from Example 3.5 are isomorphic as abstract real tori, they are non-isomorphic as *principally polarized* real tori in the sense of Mikhalkin and Zharkov [101]. In fact, there is an analogue of the Torelli theorem in this context saying that up to certain “Whitney flips,” the Jacobian as a principally polarized real torus determines the metric graph Γ ; see [36]. The map $AJ_{v_0} : \Gamma \rightarrow \Omega(\Gamma)^*/H_1(\Gamma, \mathbb{Z})$ is *harmonic* (or *balanced*) in a certain natural sense; see, e.g., [18, Theorem 4.1]. This fact is used in the paper [81], about which we will say more in Sect. 11.2. Basic properties of the tropical Abel–Jacobi map are also used in important ways in [41, 52].


4 Ranks of Divisors

4.1 Linear Systems

By analogy with algebraic curves, a divisor $D = \sum D(v)v$ on a (metric) graph is called *effective* if $D(v) \geq 0$ for all v , and we write $D \geq 0$. The complete linear series of a divisor D is defined to be

$$|D| = \{E \geq 0 \mid E \sim D\}.$$

Fig. 5 The linear system $|2v + w|$ is not equidimensional



Similarly, we write

$$R(D) = \{f \in \text{PL}(\Gamma) \mid \text{div}(f) + D \geq 0\}$$

for the set of tropical rational functions with poles along the divisor D .

As explained in [63], the complete linear series $|D|$ has the structure of a compact polyhedral complex. However, this polyhedral complex often fails to be equidimensional, as the following example shows.

Example 4.1. Consider the metric graph pictured in Fig. 5, consisting of two loops attached at a point v . Let $D = 2v + w$, where w is a point on the interior of the first loop. The complete linear system is the union of two tori. The first, which consists of divisors equivalent to D that are supported on the first loop, has dimension two, while the other, which consists of divisors equivalent to D that have positive degree on the second loop, has dimension one.


We now wish to define the rank of a divisor on a graph. As the previous example shows, the appropriate definition should not be the dimension of the linear system $|D|$, considered as a polyhedral complex.⁴ Instead, we note that a line bundle L on an algebraic curve C has rank at least r if and only if, for every collection of r points of C , there is a nonzero section of L that vanishes at those points. This motivates the following definition.

Definition 4.2. Let D be a divisor on a (metric) graph. If D is not equivalent to an effective divisor, we define its rank to be -1 . Otherwise, we define $r(D)$ to be the largest nonnegative integer r such that $|D - E| \neq \emptyset$ for all effective divisors E of degree r .

Example 4.3. Even when a linear series is equidimensional, its dimension may not be equal to the rank of the corresponding divisor. For example, consider the metric graph pictured in Fig. 6, consisting of a loop meeting a line segment in a vertex v . The linear system $|v|$ is 1-dimensional, because v is equivalent to any point on the line segment. The rank of v , however, is 0, because if w lies in the interior of the loop, then v is not equivalent to w .

⁴Another natural idea would be to try to define $r(D)$ as one less than the “dimension” of $R(D)$ considered as a semimodule over the tropical semiring \mathbf{T} consisting of $\mathbb{R} \cup \{\infty\}$ together with the operations of \min and $+$. However, this approach also faces significant difficulties. See [73] for a detailed discussion of the tropical semimodule structure on $R(D)$.

Fig. 6 The vertex v moves in a one-dimensional family, but has rank zero



Remark 4.4. There are several other notions of rank in the literature related to our setup of a line bundle on a degenerating family of curves. We mention, for example, the generalized rank functions of Katz and Zureick-Brown [80] and the algebraic rank of Caporaso [35]. The rank as defined here is sometimes referred to as the *combinatorial rank* to distinguish it from these other invariants.

As in the case of curves, we write

$$W_d^r(\Gamma) := \{D \in \text{Pic}^d(\Gamma) \mid r(D) \geq r\}.$$

We similarly define the *gonality* of a graph to be the smallest degree of a divisor of rank at least one. The *Clifford index* of the graph is

$$\text{Cliff}(\Gamma) := \min\{\deg(D) - 2r(D) \mid r(D) > \max\{0, \deg(D) - g + 1\}\}.$$

Remark 4.5. The definition of gonality above is sometimes called the *divisorial gonality*, to distinguish it from the *stable gonality*, which is the smallest degree of a harmonic morphism from a modification of the given metric graph to a tree. The divisorial gonality is always less than or equal to the stable gonality, see, e.g., [9].

4.2 Specialization

One of the key properties of the combinatorial rank is its behavior under specialization. Note that the specialization map takes effective line bundles to effective divisors. Combining this with the fact that it takes principal divisors to principal divisors, we see that, for any divisor D on C , we have

$$\text{Trop } |D| \subseteq |\text{Trop}(D)| \text{ and}$$

$$\text{trop } \mathcal{L}(D) \subseteq R(\text{Trop}(D)).$$

Combining this with the definition of rank yields the following semicontinuity result.

Specialization Theorem. [17] *Let D be a divisor on C . Then*

$$r(D) \leq r(\text{Trop}(D)).$$

Another way of stating this is that $\text{Trop}(W_d^r(C)) \subseteq W_d^r(\Gamma)$. The power of the Specialization Theorem lies in the fact that the rank of the divisor D is an algebro-geometric invariant, whereas the rank of $\text{Trop}(D)$ is a combinatorial invariant. We can therefore use techniques from each field to inform the other. For example, an immediate consequence of specialization is the following fact.

Theorem 4.6. *Let Γ be a metric graph of genus g , and let d, r be positive integers such that $g - (r + 1)(g - d + r) \geq 0$. Then $W_d^r(\Gamma) \neq \emptyset$.*

Proof. As we will see in Corollary 7.2, there exist a curve C over a discretely valued field K , and a semistable R -model \mathcal{C} of C such that the dual graph of the central fiber is isometric to Γ . A well-known theorem of Kempf and Kleiman–Laksov [84, 86] asserts that $W_d^r(C) \neq \emptyset$. It follows from Theorem 4.2 that $W_d^r(\Gamma) \neq \emptyset$ as well.

Corollary 4.7. *A metric graph of genus g has gonality at most $\lceil \frac{g+2}{2} \rceil$.*

Remark 4.8. We are unaware of a purely combinatorial proof of Theorem 4.6. There are many reasons that such a proof would be of independent interest. For example, a combinatorial proof could shed some light on whether the analogous statement is true for finite graphs, as conjectured in [17, Conjectures 3.10, 3.14].⁵

4.3 Riemann–Roch

Another key property of the combinatorial rank is that it satisfies a tropical analogue of the Riemann–Roch theorem:

Tropical Riemann–Roch Theorem ([63, 101]). *Let Γ be a metric graph of genus g , and let $K_\Gamma := \sum_{v \in \Gamma} (\text{val}(v) - 2)v$ be the canonical divisor on Γ . Then for every divisor D on Γ ,*

$$r(D) - r(K_\Gamma - D) = \text{deg}(D) - g + 1.$$

The first result of this kind was the discrete analogue of Theorem 4.3 for finite (non-metric) graphs proved in [19]. In the Baker–Norine Riemann–Roch theorem, one defines $r(D)$ for a divisor D on a graph G exactly as in Definition 4.2, the subtle difference being that the effective divisor E is restricted to the vertices of G . Gathmann and Kerber [63] showed that one can deduce Theorem 4.3 from the Baker–Norine theorem using a clever approximation argument, whereas Mikhalkin and Zharkov [101] generalized the method of proof from [19] to the metric graph

⁵Sam Payne has pointed out that there is a gap in the proof of Conjectures 3.10 and 3.14 of Baker [17] given in [34]. The claim on page 82 that $W_{d,\phi}^r$ has nonempty fiber over b_0 does not follow from the discussion that precedes it. A priori, the fiber of $\overline{W_{d,\phi}^r}$ over b_0 might be contained in the boundary $\overline{P_\phi^d} \setminus \text{Pic}_\phi^d$ of the compactified relative Picard scheme.

setting. Later on, Hladky–Kral–Norine [75] and Luo [96] proved theorems which imply that one can also deduce Riemann–Roch for graphs from tropical Riemann–Roch. (We will discuss Luo’s theory of rank-determining sets in Sect. 5.2.) So in retrospect, one can say that in some sense the Baker–Norine theorem and the theorem stated above are equivalent.

Baker and Norine’s strategy of proof for Theorem 4.3, as modified by Mikhalkin and Zharkov, is to first show that if \mathcal{O} is an orientation of the graph (i.e., a choice of a head vertex and tail vertex for each edge of G), then

$$D_{\mathcal{O}} := \sum_{v \in V(G)} (\text{indeg}_{\mathcal{O}}(v) - 1)v$$

is a divisor of degree $g - 1$, and this divisor has rank -1 if and only if the orientation \mathcal{O} is acyclic. This fact helps to establish the Riemann–Roch theorem in the case of divisors of degree $g - 1$, which serves as the base case for the more general argument. It is interesting to note that the tropical Riemann–Roch theorem has thus far resisted attempts to prove it via classical algebraic geometry. At present, neither the tropical Riemann–Roch theorem nor the classical Riemann–Roch theorem for algebraic curves is known to imply the other.

If \mathcal{C} is a strongly semistable R -model for a curve C over a discretely valued field K with the property that all irreducible components of the special fiber C_0 have genus 0, then the multidegree of the relative dualizing sheaf $\Omega_{C/R}^1$ is equal to the canonical divisor of the graph G . This is a simple consequence of the adjunction formula, which shows more generally that $\text{mdeg}(\Omega_{C/R}^1) = K_{(G,\omega)}^\#$ in the terminology of Sect. 4.4 below. If K is not discretely valued, this is still true with the right definition of the sheaf $\Omega_{C/R}^1$ (see [81]). This “explains” in some sense why there is a canonical *divisor* on a metric graph while on an algebraic curve there is merely a canonical *divisor class*.

It is clear from the definition of rank that if D and E are divisors on a metric graph Γ having nonnegative rank, then $r(D + E) \geq r(D) + r(E)$. Combining this with tropical Riemann–Roch, one obtains a tropical version of Clifford’s inequality:

Tropical Clifford’s Theorem. *Let D be a special divisor on a metric graph Γ , that is, a divisor such that both D and $K_\Gamma - D$ have nonnegative rank. Then*

$$r(D) \leq \frac{1}{2} \text{deg}(D).$$

Remark 4.9. The classical version of Clifford’s theorem is typically stated in two parts. The first part is the inequality above, while the second part states that, when equality holds, either $D \sim 0$, $D \sim K_C$ or the curve C is hyperelliptic and the linear equivalence class of D is a multiple of the unique g_2^1 . The same conditions for equality hold in the tropical case as well, by a recent theorem of Coppens, but the proof is quite subtle as the classical methods do not work in the tropical context. See [52] for details.




Fig. 7 Two metric graphs of genus 3, the first of which is hyperelliptic, and the second of which is not. (All edges have length 1)

Note that, as in the case of curves, the Riemann–Roch theorem significantly limits the possible ranks that a divisor of fixed degree on a metric graph may have. For example, a divisor of negative degree necessarily has rank -1 , so a divisor of degree $d > 2g - 2$ must have rank $d - g$. It is only in the intermediate range $0 \leq d \leq 2g - 2$ where there are multiple possibilities for the rank.

Example 4.10. The canonical divisor on a circle is trivial, and it is the only divisor of degree 0 with nonnegative rank. If D is a divisor of degree $d > 0$, then by Riemann–Roch D has rank $d - 1$. This can also be seen using the fact that the circle is a torsor for its Jacobian, as in Example 3.4: if E is an effective divisor of degree $d - 1$, then there is a unique point P such that $D - E \sim P$, and hence $r(D) \geq d - 1$.

Example 4.11. The smallest genus for which the rank of an effective divisor is not completely determined by the degree is genus 3. Pictured in Fig. 7 are two examples of genus 3 metric graphs, the first of which is *hyperelliptic*, meaning that it admits a divisor of degree 2 and rank 1, and the second of which is not. For the first graph, one can check by hand that the sum of the two vertices on the left has rank at least 1, and it cannot have rank higher than 1 by Clifford’s theorem. We will show that the second graph is not hyperelliptic in Example 5.4, as the argument will require some techniques for computing ranks of divisors that we will discuss in the next section.

4.4 Divisors on Vertex-Weighted Graphs

In [6], Amini and Caporaso formulate a refinement of the Specialization Theorem which takes into account the genera of the components of the special fiber. In this section we describe their result following the presentation in [5].

A *vertex-weighted metric graph* is a pair (Γ, ω) consisting of a metric graph Γ and a weight function $\omega : \Gamma \rightarrow \mathbb{Z}_{\geq 0}$ such that $\omega(x) = 0$ for all but finitely many $x \in \Gamma$. Following [6], we define a new metric graph $\Gamma^\#$ by attaching $\omega(x)$ loops of arbitrary positive length at each point $x \in \Gamma$. There is a natural inclusion of Γ into $\Gamma^\#$. The *canonical divisor* of (Γ, ω) is defined to be

$$K^\# = K_\Gamma + \sum_{x \in \Gamma} 2\omega(x),$$

which can naturally be identified with the canonical divisor of $K_{\Gamma^\#}$ restricted to Γ . Its degree is $2g^\# - 2$, where $g^\# = g(\Gamma) + \sum_{x \in \Gamma} \omega(x)$ is the genus of $\Gamma^\#$.

Following [6], the *weighted rank* $r^\#$ of a divisor D on Γ is defined to be $r^\#(D) := r_{\Gamma^\#}(D)$. By [5, Corollary 4.12], we have the more intrinsic description

$$r^\#(D) = \min_{0 \leq E \leq \mathcal{W}} (\deg(E) + r_\Gamma(D - 2E)),$$

where $\mathcal{W} = \sum_{x \in \Gamma} \omega(x)(x)$.

The Riemann–Roch theorem for $\Gamma^\#$ implies the following “vertex-weighted” Riemann–Roch theorem for Γ :

$$r^\#(D) - r^\#(K^\# - D) = \deg(D) + 1 - g^\#.$$

If C is a curve over K , together with a semistable model \mathcal{C} over R , we define the associated vertex-weighted metric graph (Γ, ω) by taking Γ to be the skeleton of \mathcal{C} and defining the weight function ω by $\omega(v) = g_v$. With this definition, the genus of the weighted metric graph (Γ, ω) is equal to the genus of C and we have $\text{trop}(K_C) \sim K^\#$ on Γ [5, Sect. 4.7.1]. The following weighted version of the Specialization Theorem, inspired by the results of [6], is proved in [5, Theorem 4.13]:

Weighted Specialization Theorem. *For every divisor $D \in \text{Div}(C)$, we have $r_C(D) \leq r^\#(\text{trop}(D))$.*

5 Combinatorial Techniques

The tropical approach to degeneration of line bundles in algebraic geometry derives its power from the combinatorial tools which one has available, many of which have no classical analogues. We describe some of these tools in this section.

5.1 Reduced Divisors and Dhar’s Burning Algorithm

Definition 5.1. Let G be a finite graph. Given a divisor $D \in \text{Div}(G)$ and a vertex v of G , we say that D is *v-reduced* if

- (RD1) $D(w) \geq 0$ for every $w \neq v$ and
- (RD2) for every nonempty set $A \subseteq V(G) \setminus \{v\}$ there is a vertex $w \in A$ such that $\text{outdeg}_A(w) > D(w)$.

Here $\text{outdeg}_A(w)$ denotes the outdegree of w with respect to A , i.e., the number of edges connecting $w \in A$ to a vertex not in A . The following important result (cf. [19, Proposition 3.1]) shows that v -reduced divisors form a distinguished set of representatives for linear equivalence classes of divisors on G :

Lemma 5.2. *Every divisor on G is equivalent to a unique v -reduced divisor.*

If D has nonnegative rank, the v -reduced divisor equivalent to D is the divisor in $|D|$ that is lexicographically “closest” to v . It is a discrete analogue of the unique divisor in a classical linear series $|D|$ with the highest possible order of vanishing at a given point $p \in C$.

There is a simple algorithm for determining whether a given divisor satisfying (RD1) above is v -reduced, known as *Dhar’s burning algorithm*. For $w \neq v$, imagine that there are $D(w)$ buckets of water at w . Now, light a fire at v . The fire starts spreading through the graph, burning through an edge as soon as one of its endpoints is burnt, and burning a vertex w if the number of burnt edges adjacent to w is greater than $D(w)$ (that is, there is not enough water to fight the fire). The divisor D is v -reduced if and only if the fire consumes the whole graph. For a detailed account of this algorithm, we refer to [55] and the more recent [22, Sect. 5.1].

There is a completely analogous set of definitions and results for metric graphs. Let Γ be a metric graph, let D be a divisor on Γ , and choose a point $v \in \Gamma$ (which need not be a vertex). We say that D is v -reduced if the two conditions from Definition 5.1 hold, with the second condition replaced by

(RD2’) for every closed, connected, nonempty set $A \subseteq \Gamma \setminus \{v\}$ there is a point $w \in A$ such that $\text{outdeg}_A(w) > D(w)$.

It is not hard to see that condition (RD2’) is equivalent to requiring that for every non-constant tropical rational function $f \in R(\Gamma)$ with a global maximum at v , the divisor $D' := D + \text{div}(f)$ does not satisfy (RD1), i.e., there exists $w \neq v$ in Γ such that $D'(w) < 0$.

The analogue of Lemma 5.2 remains true in the metric graph context: every divisor on Γ is equivalent to a unique v -reduced divisor. Moreover, Dhar’s burning algorithm as formulated above holds almost *verbatim* for metric graphs: the fire starts spreading through Γ , getting blocked at a point $w \in \Gamma$ iff the number of burnt tangent directions at w is less than or equal to $D(w)$; the divisor D is v -reduced if and only if the fire consumes all of Γ .

We note the following important fact, which is useful for computing ranks of divisors.

Lemma 5.3. *Let D be a divisor on a finite or metric graph, and suppose that D is v -reduced for some v . If D has nonnegative rank, then $D(v) \geq r(D)$.*

Example 5.4. Let Γ be the complete graph on 4 vertices endowed with arbitrary edge lengths. We can use the theory of reduced divisors to show that Γ is not hyperelliptic, justifying one of the claims in Example 4.11. Suppose that there exists a divisor D on Γ of degree 2 and rank 1 and choose a vertex v . Since D has rank 1, $D \sim D' := v + v'$ for some $v' \in \Gamma$. Now, let $w \neq v, v'$ be a vertex. Note that there




Fig. 8 Using Dhar’s burning algorithm to compute the v_5 -reduced divisor equivalent to $v_1 + v_2$. The burnt vertices after each iteration are colored *white*

are at least two paths from w to v that do not pass through v' , and if $v = v'$, there are three. It follows by Dhar’s burning algorithm that D' is w -reduced. But $D'(w) = 0$, contradicting the fact that D (and hence D') has rank 1.

If a given divisor is not v -reduced, Dhar’s burning algorithm provides a method for finding the unique equivalent v -reduced divisor. In the case of finite graphs, after performing Dhar’s burning algorithm, if we fire the vertices that are left unburnt we obtain a divisor that is lexicographically closer to the v -reduced divisor, and after iterating the procedure a finite number of times, it terminates with the v -reduced divisor (cf. [22]). For metric graphs, there is a similar procedure but with additional subtleties—we refer the interested reader to [96, Algorithm 2.5] and [14].

Example 5.5. Consider the finite graph depicted in Fig. 8, consisting of two triangles meeting at a vertex v_3 . We let $D = v_1 + v_2$, and compute the v_5 -reduced divisor equivalent to D . After performing Dhar’s burning algorithm once, we see that vertices v_1 and v_2 are left unburnt. Firing these, we see that D is equivalent to $2v_3$. Performing Dhar’s burning algorithm a second time, all three of the vertices v_1, v_2, v_3 are not burnt. Firing these, we obtain the divisor $v_4 + v_5$. A third run of Dhar’s burning algorithm shows that $v_4 + v_5$ is v_5 -reduced.

5.2 Rank-Determining Sets

The definition of the rank of a divisor on a finite graph G implies easily that there is an algorithm for computing it.⁶ Indeed, since there are only a finite number of effective divisors E of a given degree on G , we are reduced to the problem of determining whether a given divisor is equivalent to an effective divisor or not. This problem can be solved in polynomial time by using the iterated version of Dhar’s algorithm described above to compute the v -reduced divisor D' equivalent to D for some vertex v . If $D'(v) < 0$, then $|D| = \emptyset$, and otherwise $|D| \neq \emptyset$.

⁶Although there is no known *efficient* algorithm; indeed, it is proved in [85] that this problem is NP-hard.

If one attempts to generalize this algorithm to the case of metric graphs, there is an immediate problem, since there are now an infinite number of effective divisors E to test. The idea behind rank-determining sets is that it suffices, in the definition of $r(D)$, to restrict to a finite set of effective divisors E .

Definition 5.6. Let Γ be a metric graph. A subset $A \subseteq \Gamma$ is a *rank-determining set* if, for any divisor D on Γ , D has rank at least r if and only if $|D - E| \neq \emptyset$ for every effective divisor E of degree r supported on A .

In [96], Luo provides a criterion for a subset of a metric graph to be rank-determining. A different proof of Luo’s criterion has been given recently by Backman [15]. Luo defines a *special open set* to be a connected open set $U \subseteq \Gamma$ such that every connected component $X \subseteq \Gamma \setminus U$ contains a boundary point v such that $\text{outdeg}_X(v) \geq 2$.

Theorem 5.7 ([96]). A subset $A \subseteq \Gamma$ is rank-determining if and only if all nonempty special open subsets of Γ intersect A .

Corollary 5.8. Let G be a model for a metric graph Γ . If G has no loops, then the vertices of G are a rank-determining set.

There are lots of other interesting rank-determining sets besides vertices of models.

Example 5.9. Let G be a model for a metric graph Γ . Choose a spanning tree of G and let e_1, \dots, e_g be the edges in the complement of the spanning tree. For each such edge e_i , choose a point v_i in its interior, and let w be any other point of Γ . Then the set $A = \{v_1, \dots, v_g\} \cup \{w\}$ is rank-determining. This construction is used, for example, in [52].

Example 5.10. Let G be a bipartite finite graph, and let Γ be a metric graph having G as a model. If we fix a 2-coloring of the vertices of G , then the vertices of one color are a rank-determining set. This is the key observation in [77], in which the second author shows that the Heawood graph admits a divisor of degree 7 and rank 2, regardless of the choice of edge lengths. The interest in this example arises because it shows that there is a nonempty open subset of the (highest-dimensional component of the) moduli space M_8^{trop} containing no Brill–Noether general metric graph. (See Sect. 7.1 for a description of M_g^{trop} .)

5.3 Tropical Independence

Many interesting questions about algebraic curves concern the ranks of linear maps between the vector spaces $\mathcal{L}(D)$. For example, both the Gieseker–Petri theorem and the Maximal Rank Conjecture are statements about the rank of the multiplication maps

$$\mu : \mathcal{L}(D) \otimes \mathcal{L}(D') \rightarrow \mathcal{L}(D + D')$$

for certain pairs of divisors D, D' on a general curve.

One simple strategy for showing that a map, such as μ , has rank at least k is to carefully choose k elements of the image, and then check that they are linearly independent. To this end, we formulate a notion of *tropical independence*, which gives a sufficient condition for linear independence of rational functions on a curve C in terms of the associated piecewise linear functions on the metric graph Γ .

Definition 5.11 ([78]). A set of piecewise linear functions $\{f_1, \dots, f_k\}$ on a metric graph Γ is *tropically dependent* if there are real numbers b_1, \dots, b_k such that the minimum

$$\min\{f_1(v) + b_1, \dots, f_k(v) + b_k\}$$

occurs at least twice at every point v in Γ .

If there are no such real numbers b_1, \dots, b_k , then we say $\{f_1, \dots, f_k\}$ is *tropically independent*. We note that linearly dependent functions on C specialize to tropically dependent functions on Γ . Although the definition of tropical independence is merely a translation of linear dependence into tropical language, one can often check tropical independence using combinatorial methods. The following lemma illustrates this idea.

Lemma 5.12 ([78]). Let D be a divisor on a metric graph Γ , with f_1, \dots, f_k piecewise linear functions in $R(D)$, and let

$$\theta = \min\{f_1, \dots, f_k\}.$$

Let $\Gamma_j \subset \Gamma$ be the closed set where $\theta = f_j$, and let $v \in \Gamma_j$. Then the support of $\text{div}(\theta) + D$ contains v if and only if v belongs to either

1. the support of $\text{div}(f_j) + D$ or
2. the boundary of Γ_j .

5.4 Break Divisors

Another useful combinatorial tool for studying divisor classes on graphs and metric graphs is provided by the theory of *break divisors*, which was initiated by Mikhalkin–Zharkov in [101] and studied further by An–Baker–Kuperberg–Shokrieh in [10]. Given a metric graph Γ of genus g , fix a model G for Γ . For each spanning tree T of G , let Σ_T be the image of the canonical map

$$\prod_{e \notin T} e \rightarrow \text{Div}_+^g(\Gamma)$$

sending (p_1, \dots, p_g) to $p_1 + \dots + p_g$. (Here $\text{Div}_+^g(\Gamma)$ denotes the set of effective divisors of degree g on Γ and e denotes a *closed* edge of G , so the points p_i are allowed to be vertices of G .) We call $B(\Gamma) := \bigcup_T \Sigma_T$ the set of *break divisors* on Γ . The set of break divisors does not depend on the choice of the model G . The following result shows that the natural map $B(\Gamma) \subset \text{Div}_+^g(\Gamma) \rightarrow \text{Pic}^g(\Gamma)$ is bijective:

Theorem 5.13 ([10, 101]). *Every divisor of degree g on Γ is linearly equivalent to a unique break divisor.*

Since $B(\Gamma)$ is a compact subset of $\text{Div}_+^g(\Gamma)$ and $\text{Pic}^g(\Gamma)$ is also compact, it follows from general topology that there is a canonical continuous section σ to the natural map $\pi : \text{Div}_+^g(\Gamma) \rightarrow \text{Pic}^g(\Gamma)$ whose image is precisely the set of break divisors. In particular, every degree g divisor class on a metric graph Γ has a *canonical* effective representative. The analogue of this statement in algebraic geometry is *false*: when $g = 2$, for example, the natural map $\text{Sym}^2(C) = \text{Div}_+^2(C) \rightarrow \text{Pic}^2(C)$ is a birational isomorphism which blows down the \mathbb{P}^1 corresponding to the fiber over the unique \mathfrak{g}_2^1 , and this map has no section. This highlights an interesting difference between the algebraic and tropical settings.


The proof of Theorem 5.13 in [101] utilizes the theory of tropical theta functions and the tropical analogue of Riemann’s theta constant. A purely combinatorial proof based on the theory of *q-connected orientations* is given in [10], and the combinatorial proof yields an interesting analogue of Theorem 5.13 for finite graphs. If G is a finite graph and Γ is its *regular realization*, in which all edges are assigned a length of 1, define the set of *integral break divisors* on G to be $B(G) = B(\Gamma) \cap \text{Div}(G)$. In other words, $B(G)$ consists of all break divisors for the underlying metric graph Γ which are supported on the vertices of G .

Theorem 5.14 ([10]). *Every divisor of degree g on G is linearly equivalent to a unique integral break divisor.*

Since $\text{Pic}^g(G)$ and $\text{Pic}^0(G) = \text{Jac}(G)$ have the same cardinality (the former is naturally a torsor for the latter), it follows from Remark 2.4 that the number of integral break divisors on G is equal to the number of spanning trees of G , though there is in general no canonical bijection between the two. A family of interesting combinatorial bijections is discussed in [23].

If we define $C_T = \pi(\Sigma_T)$, then $\text{Pic}^g(\Gamma) = \bigcup_T C_T$ by Theorem 5.13. It turns out that the relative interior of each cell C_T is (the interior of) a parallelotope, and if $T \neq T'$, then the relative interiors of C_T and $C_{T'}$ are disjoint. Thus $\text{Pic}^g(\Gamma)$ has a polyhedral decomposition depending only on the choice of a model for Γ . The maximal cells in the decomposition correspond naturally to spanning trees, and the minimal cells (i.e., vertices) correspond naturally to integral break divisors, as illustrated in Fig. 9.

Fig. 9 A polyhedral decomposition of $\text{Pic}^2(\Gamma)$ for the metric realization of the graph G obtained by deleting an edge from the complete graph K_4



Remark 5.15. Mumford’s non-Archimedean analytic uniformization theory for degenerating abelian varieties [102], as recently refined by Gubler and translated into the language of tropical geometry [67, 68], shows that if G is the dual graph of the special fiber of a regular R -model \mathcal{C} for a curve C , then the canonical polyhedral decomposition $\{C_T\}$ of $\text{Pic}^g(\Gamma)$ gives rise to a canonical proper R -model for $\text{Pic}^g(C)$. Sam Payne has asked (Payne, 2013, Personal communication) whether (up to the identification of Pic^g with Pic^0) this model coincides with the compactification of the Néron model of $\text{Jac}(C)$ introduced by Caporaso in [32].

Break divisors corresponding to the relative interior of some cell C_T are called *simple break divisors*. They can be characterized as the set of degree g effective divisors D on Γ such that $\Gamma \setminus \text{supp}(D)$ is connected and simply connected. Dhar’s algorithm shows that such divisors are *universally reduced*, i.e., they are q -reduced for all $q \in \Gamma$. A consequence of this observation and the Riemann–Roch theorem for metric graphs is the following result, which is useful in tropical Brill–Noether theory (cf. Sect. 9.2).

Proposition 5.16. *Let Γ be a metric graph and let D be a simple break divisor (or more generally any universally reduced divisor) on Γ . Then D has rank 0 and $K_\Gamma - D$ has rank -1 . Therefore:*

- (1) *The set of divisor classes in $\text{Pic}^g(\Gamma)$ having rank at least 1 is contained in the codimension one skeleton of the polyhedral decomposition $\bigcup_T C_T$.*
- (2) *If T is a spanning tree for some model G of Γ and D, E are effective divisors with $D + E$ linearly equivalent to K_Γ , then there must be an open edge e° in the complement of T such that D has no chips on e° .*

Remark 5.17. The set $B(\Gamma)$ of all break divisors on Γ can be characterized as the topological closure in $\text{Div}_+^g(\Gamma)$ of the set of universally reduced divisors.

Remark 5.18. The real torus $\text{Pic}^g(\Gamma)$ has a natural Riemannian metric. One can compute the volume of $\text{Pic}^g(\Gamma)$ in terms of a matrix determinant associated with G , and the volume of the cell C_T is the product of the lengths of the edges not in T . Comparing the volume of the torus to the sums of the volumes of the cells C_T yields a *dual version* of a weighted form of Kirchhoff’s Matrix-Tree Theorem. See [10] for details.

Part 2: Advanced Topics

We now turn to the more advanced topics of non-Archimedean analysis, tropical moduli spaces, and metrized complexes. Each of these topics plays an important role in tropical Brill–Noether theory, and we would be remiss not to mention them here. We note, however, that most of the applications we discuss in Part 3 do not require these techniques, and the casual reader may wish to skip this part on the first pass.

6 Berkovich Analytic Theory

Rather than considering a curve over a discretely valued field and then examining its behavior under base change, we could instead start with a curve over an algebraically closed field and directly associate a metric graph to it. We do this by making use of Berkovich’s theory of analytic spaces. In addition to being a convenient bookkeeping device for changes in dual graphs and specialization maps under field extensions, Berkovich’s theory also allows for clean formulations of some essential results in the theory of tropical linear series, such as the Slope Formula (Theorem 6.4 below). The theory also furnishes a wealth of powerful tools for understanding the relationship between algebraic curves and their tropicalizations.

6.1 A Quick Introduction to Berkovich Spaces

We let K be a field which is complete with respect to a non-Archimedean valuation

$$\text{val} : K^* \rightarrow \mathbb{R}.$$

We let $R \subset K$ be the valuation ring, κ the residue field, and $|\cdot| = \exp(-\text{val})$ the corresponding norm on K .

A *multiplicative seminorm* on a nonzero ring A is a function $|\cdot| : A \rightarrow \mathbb{R}$ such that for all $x, y \in A$ we have $|0| = 0$, $|1| = 1$, $|xy| = |x| \cdot |y|$, and $|x + y| \leq |x| + |y|$. A multiplicative seminorm is *non-Archimedean* if $|x + y| \leq \max\{|x|, |y|\}$ for all $x, y \in A$, and is a *norm* if $|x| = 0$ implies $x = 0$. If L is a field, a function $|\cdot| : L \rightarrow \mathbb{R}$ is a non-Archimedean norm if and only if $-\log |\cdot| : L \rightarrow \mathbb{R} \cup \{\infty\}$ is a *valuation* in the sense of Krull.

If $X = \text{Spec}(A)$ is an affine scheme over K , we define its *Berkovich analytification* X^{an} to be the set of non-Archimedean multiplicative seminorms on the K -algebra A extending the given absolute value on K , endowed with the weakest topology such that the map $X^{\text{an}} \rightarrow \mathbb{R}$ defined by $|\cdot|_x \mapsto |f|_x$ is continuous for all $f \in A$. One can globalize this construction to give a Berkovich analytification of an arbitrary scheme of finite type over K . As we have defined it, the Berkovich analytification is merely a topological space, but it can be equipped with a structure similar to that of a locally ringed space and one can view X^{an} as an object in a larger category of (not necessarily algebraizable) *Berkovich analytic spaces*. The space X^{an} is locally compact and locally path-connected. It is Hausdorff if and only if X is separated, compact if and only if X is proper, and path-connected if and only if X is connected. We refer the reader to [29, 51] for more background information on Berkovich spaces in general, and [16, 25] for more details in the special case of curves.

There is an alternate perspective on Berkovich spaces which is often useful and highlights the close analogy with schemes. If K is a field, points of an affine K -scheme $\text{Spec}(A)$ can be identified with equivalence classes of pairs (L, ϕ) where L is a field extension of K , $\phi : A \rightarrow L$ is a K -algebra homomorphism, and two pairs (L_1, ϕ_1) and (L_2, ϕ_2) are equivalent if there are embeddings of L_1 and L_2 into a common overfield L' and a homomorphism $\phi' : A \rightarrow L'$ such that the composition $\phi_i : A \rightarrow L_i \rightarrow L$ is ϕ' . Indeed, to a pair (L, ϕ) one can associate the prime ideal $\ker(\phi)$ of A , and to a prime ideal \mathfrak{p} of A one can associate the pair $(K(\mathfrak{p}), \phi)$ where $K(\mathfrak{p})$ is the fraction field of A/\mathfrak{p} and $\phi : A \rightarrow K(\mathfrak{p})$ is the canonical map.

Similarly, if K is a complete valued field, points of $\text{Spec}(A)^{\text{an}}$ can be identified with equivalence classes of pairs (L, ϕ) , where L is a complete valued field extension of K and $\phi : A \rightarrow L$ is a K -algebra homomorphism. The equivalence relation is as before, except that L' should be complete and extend the valuation on the L_i . Indeed, to a pair (L, ϕ) one can associate the multiplicative seminorm $a \mapsto |\phi(a)|$ on A , and to a multiplicative seminorm $|\cdot|_x$ one can associate the pair $(\mathcal{H}(x), \phi)$ where $\mathcal{H}(x)$ is the completion of the fraction field of $A/\ker(|\cdot|_x)$ and $\phi : A \rightarrow \mathcal{H}(x)$ is the canonical map.

We will assume for the rest of this section that K is **algebraically closed and nontrivially valued**. This ensures, for example, that the set $X(K)$ is dense in X^{an} .

If X/K is an irreducible variety, there is a dense subset of X^{an} consisting of the set Val_X of norms⁷ on the function field $K(X)$ that extend the given norm on K . Within the set Val_X , there is a distinguished class of norms corresponding to *divisorial valuations*. By definition, a valuation v on $K(X)$ is *divisorial* if there is an R -model \mathcal{X} for X and an irreducible component Z of the special fiber of \mathcal{X} such that $v(f)$ is equal to the order of vanishing of f along Z . The set of divisorial points is known to be dense in X^{an} .

Remark 6.1. In this survey we have intentionally avoided the traditional perspective of tropical geometry, in which one considers subvarieties of the torus $(K^*)^n$, and the tropicalization is simply the image of coordinatewise valuation. We refer the reader to [98] for a detailed account of this viewpoint on tropical geometry. The Berkovich analytification can be thought of as a sort of intrinsic tropicalization—the one that does not depend on a choice of coordinates. This is reinforced by the result that the Berkovich analytification is the inverse limit of all tropicalizations [62, 108].

6.2 Berkovich Curves and Their Skeleta

If C/K is a complete nonsingular curve, the underlying set of the Berkovich analytic space C^{an} consists of the points of $C(K)$ together with the set Val_C . We write

$$\text{val}_y : K(C)^* \rightarrow \mathbb{R}$$

for the valuation corresponding to a point $y \in \text{Val}_C = C^{\text{an}} \setminus C(K)$. The points in $C(K)$ are called type-1 points, and the remaining points of C^{an} are classified into three more types. We will not define points of type 3 or 4 in this survey article; see, e.g., [24, Sect. 3.5] for a definition. Note, however, that every point of C^{an} becomes a type-1 point after base-changing to a suitably large complete non-Archimedean field extension L/K .

If the residue field of $K(C)$ with respect to val_y has transcendence degree 1 over κ , then y is called a *type-2 point*. These are exactly the points corresponding to divisorial valuations. Because it has transcendence degree 1 over κ , this residue field corresponds to a unique smooth projective curve over κ , which we denote C_y . A *tangent direction* at y is an equivalence class of continuous injections $\gamma : [0, 1] \rightarrow C^{\text{an}}$ sending 0 to y , where $\gamma \sim \gamma'$ if $\gamma([0, 1]) \cap \gamma'([0, 1]) \not\supseteq \{y\}$. Closed points of C_y are in one-to-one correspondence with tangent directions at y in C^{an} .

There is natural metric on the set Val_C which is described in detail in [25, Sect. 5.3]. This metric induces a topology that is much finer than the subspace

⁷Note that there is a one-to-one correspondence between norms on $K(X)$ and valuations on $K(X)$, hence the terminology Val_X . It is often convenient to work with (semi-)valuations rather than (semi-)norms.

topology on $\text{Val}_C \subset C^{\text{an}}$, and with respect to this metric, Val_C is locally an \mathbb{R} -tree⁸ with branching precisely at the type-2 points. The type-1 points should be thought of as infinitely far away from every point of Val_C .

The local \mathbb{R} -tree structure arises in the following way. If \mathcal{C} is an R -model for C and Z is a reduced and irreducible component of the special fiber of \mathcal{C} , then Z corresponds to a type-2 point y_Z of \mathcal{C} . Blowing up a nonsingular closed point of Z (with respect to some choice of a uniformizer $\varpi \in \mathfrak{m}_R$) gives a new point $y_{Z'}$ of C^{an} corresponding to the exceptional divisor Z' of the blowup. We can then blow up a nonsingular closed point on the exceptional divisor Z' to obtain a new point of C^{an} , and so forth. The resulting constellation of points obtained by all such sequences of blowups, and varying over all possible choices of ϖ , is an \mathbb{R} -tree T_Z rooted at y_Z , as pictured in Fig. 10. The distance between the points y_Z and $y_{Z'}$ is $\text{val}(\varpi)$.

A *semistable vertex set* is a finite set of type-2 points whose complement is a disjoint union of finitely many open annuli and infinitely many open balls. There is a one-to-one correspondence between semistable vertex sets and semistable models of C . More specifically, the normalized components of the central fiber of this semistable model are precisely the curves C_y for y in the semistable vertex set, and the preimages of the nodes under specialization are the annuli. The annulus corresponding to a node where C_y meets $C_{y'}$ contains a unique open segment with endpoints y and y' , and its length (with respect to the natural metric on Val_C) is the logarithmic modulus of the annulus. The union of these open segments together with the semistable vertex set is a closed connected metric graph Γ contained in C^{an} , called the *skeleton* of the semistable model \mathcal{C} . If C has genus at least 2, then there is a unique minimal semistable vertex set in C^{an} and a corresponding minimal skeleton.

Fix a semistable model \mathcal{C} of C and a corresponding skeleton $\Gamma = \Gamma_{\mathcal{C}}$. Each connected component of $C^{\text{an}} \setminus \Gamma$ has a unique boundary point in Γ , and there is a canonical retraction to the skeleton

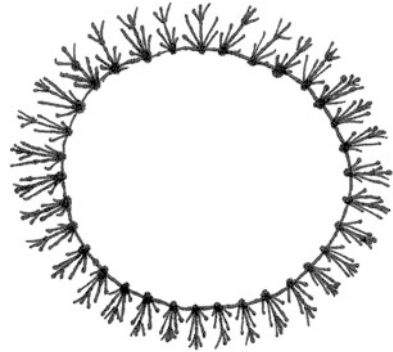
$$\tau : C^{\text{an}} \rightarrow \Gamma$$

taking a connected component of $C^{\text{an}} \setminus \Gamma$ to its boundary point. There is a natural homeomorphism of topological spaces $C^{\text{an}} \cong \varprojlim \Gamma_{\mathcal{C}}$, where the inverse limit is taken over all semistable models \mathcal{C} (cf. [25, Theorem 5.2]).

Example 6.2. Figure 10 depicts the Berkovich analytification of an elliptic curve E/K with non-integral j -invariant j_E . In this case, the skeleton Γ associated with a minimal proper semistable model of R is isometric to a circle with circumference $-\text{val}(j_E)$. There are an infinite number of infinitely branched \mathbb{R} -trees emanating from the circle at each type-2 point of the skeleton. The retraction map takes a point

⁸See [21, Appendix B] for an introduction to the theory of \mathbb{R} -trees. For our purposes, what is most important about \mathbb{R} -trees is that there is a unique path between any two points.

Fig. 10 The skeleton of an elliptic curve with non-integral j -invariant



$x \in E^{\text{an}}$ to the endpoint in Γ of the unique path from x to Γ . The points of $E(K)$ lie “out at infinity” in the picture: they are ends of the \mathbb{R} -trees.

6.3 Tropicalization of Divisors and Functions on Curves

Restricting the retraction map τ to $C(K)$ and extending linearly give the tropicalization map on divisors

$$\text{Trop} : \text{Div}(C) \rightarrow \text{Div}(\Gamma).$$

If K_0 is a discretely valued subfield of K over which C is defined and has semistable reduction, and if D is a divisor on C whose support consists of K_0 -rational points, then the divisor $\text{mdeg}(D)$ on G (identified with a divisor on Γ via the natural inclusion) coincides with the tropicalization $\text{Trop}(D)$ defined via retraction to the skeleton.

Example 6.3. Returning to Example 3.1, in which the metric graph Γ is a closed line segment of length 1, by considering the divisor cut out by $y^a = x^b z^{a-b}$ for positive integers $a > b$ we see that a divisor on $C_{\bar{K}}$ can tropicalize to any rational point on Γ .

Given a rational function $f \in K(C)^*$, we write $\text{trop}(f)$ for the real valued function on the skeleton Γ given by $y \mapsto \text{val}_y(f)$. The function $\text{trop}(f)$ is piecewise linear with integer slopes, and thus we obtain a map

$$\text{trop} : K(C)^* \rightarrow \text{PL}(\Gamma).$$

Moreover, this map respects linear equivalence of divisors, in the sense that if $D \sim D'$ on C then $\text{trop}(D) \sim \text{trop}(D')$ on Γ . In particular, the tropicalization map on divisors descends to a natural map on Picard groups

$$\text{Trop} : \text{Pic}(C) \rightarrow \text{Pic}(\Gamma).$$

One can refine this observation as follows. Let x be a type-2 point in C^{an} . Given a nonzero rational function f on C , one can define its *normalized reduction* \tilde{f}_x with respect to x as follows. Choose $c \in K^*$ such that $|f|_x = |c|$. Define $\tilde{f}_x \in \kappa(C_x)^*$ to be the image of $c^{-1}f$ in the residue field of $K(C)$ with respect to val_x , which by definition is isomorphic to $\kappa(C_x)$. Although f_x is only well-defined up to multiplication by an element of κ^* , its divisor $\text{div}(f_x)$ is completely well-defined. We define the normalized reduction of the zero function to be zero.

Given an $(r + 1)$ -dimensional K -vector space $H \subset K(C)$, its reduction $\tilde{H}_x = \{\tilde{f}_x \mid f \in H\}$ is an $(r + 1)$ -dimensional vector space over κ . Given \tilde{f} in the function field of C_x and a closed point ν of C_x , we let $s^\nu(\tilde{f}) := \text{ord}_\nu(\tilde{f})$ be the order of vanishing of \tilde{f} at ν . If $\tilde{f} = \tilde{f}_x$ for $f \in K(C)^*$, then $s^\nu(\tilde{f})$ is equal to the slope of $\text{trop}(f)$ in the tangent direction at x corresponding to ν . This is a consequence of the non-Archimedean Poincaré–Lelong formula, due to Thuillier [116]. Using this observation, one deduces the following important result (cf. [25, Theorem 5.15]):

Theorem 6.4 (Slope Formula). *For any nonzero rational function $f \in K(C)$,*

$$\text{Trop}(\text{div}(f)) = \text{div}(\text{trop}(f)).$$

6.4 Skeletons of Higher-Dimensional Berkovich Spaces

The construction of the skeleton of a semistable model of a curve given in Sect. 6.2 can be generalized in various ways to higher dimensions. For brevity, we mention just three such generalizations. In what follows, X will denote a proper variety of dimension n over K .

1. *Semistable models.* Suppose \mathcal{X} is a strictly semistable model of X over R . Then the geometric realization of the dual complex $\Delta(\overline{\mathcal{X}})$ of the special fiber embeds naturally in the Berkovich analytic space X^{an} , and as in the case of curves there is a strong deformation retraction of X^{an} onto $\Delta(\overline{\mathcal{X}})$. These facts are special cases of results due to Berkovich; see [104] for a lucid explanation of the constructions in the special case of strictly semistable models, and [70] for a generalization to “extended skeleta.”
2. *Toroidal embeddings.* A *toroidal embedding* is, roughly speaking, something which looks étale-locally like a toric variety together with its dense big open torus. When K is trivially valued, Thuillier [117] associates a skeleton $\Sigma(X)$ of U and an extended skeleton $\overline{\Sigma}(X)$ of X to any toroidal embedding $U \subset X$ (see also [1]). As in the case of semistable models, the skeleton $\overline{\Sigma}(X)$ embeds naturally into X^{an} and there is a strong deformation retract $X^{\text{an}} \rightarrow \overline{\Sigma}(X)$.
3. *Abelian varieties.* If E/K is an elliptic curve with non-integral j -invariant, the skeleton associated with a minimal proper semistable model of R can also be constructed using Tate’s non-Archimedean uniformization theory. In this case,

the skeleton of E^{an} is the quotient of the skeleton of \mathbf{G}_m (which is isomorphic to \mathbb{R} and consists of the unique path from 0 to ∞ in $(\mathbb{P}^1)^{\text{an}}$) by the map $x \mapsto x - \text{val}(j_E)$. Using Mumford’s higher-dimensional generalization of Tate’s theory [102], one can define a skeleton associated with any totally degenerate abelian variety A ; it is a real torus of dimension $\dim(A)$. This can be generalized further using Raynaud’s uniformization theory to define a canonical notion of skeleton for an arbitrary abelian variety A/K (see, e.g., [68]). If A is principally polarized, there is an induced *tropical principal polarization* on the skeleton of A , see [20, Sect. 3.7] for a definition. It is shown in [20] that the skeleton of the Jacobian of a curve C is isomorphic to the Jacobian of the skeleton as principally polarized real tori:

Theorem 6.5 ([20]). *Let C be a curve over an algebraically closed field, complete with respect to a nontrivial valuation, such that the minimal skeleton of the Berkovich analytic space C^{an} is isometric to Γ . Then there is a canonical isomorphism of principally polarized real tori $\text{Jac}(\Gamma) \cong \Sigma(\text{Jac}(C)^{\text{an}})$ making the following diagram commute.*

$$\begin{array}{ccc}
 C^{\text{an}} & \xrightarrow{AJ} & \text{Jac}(C)^{\text{an}} \\
 \downarrow & & \searrow \\
 \Gamma & \xrightarrow{AJ} & \text{Jac}(\Gamma) \xrightarrow{\sim} \Sigma(\text{Jac}(C)^{\text{an}}).
 \end{array}$$

Remark 6.6. Theorem 6.5 has the following interpretation in terms of tropical moduli spaces, which we discuss in greater detail in Sect. 7. There is a map

$$\text{trop} : M_g \rightarrow M_g^{\text{trop}}$$

from the moduli space of genus $g \geq 2$ curves to the moduli space of tropical curves of genus g which takes a curve C to its minimal skeleton, considered as a vertex-weighted metric graph. There is also a map (of sets, for example)

$$\text{trop} : A_g \rightarrow A_g^{\text{trop}}$$

from the moduli space of principally polarized abelian varieties of dimension g to the moduli space of “principally polarized tropical abelian varieties” of dimension g , taking an abelian variety to its skeleton in the sense of Berkovich. Finally, there are Torelli maps $M_g \rightarrow A_g$ (resp., $M_g^{\text{trop}} \rightarrow A_g^{\text{trop}}$) which take a curve (resp., metric graph) to its Jacobian [31]. Theorem 6.5 implies that the following square commutes:

$$\begin{array}{ccc}
 M_g & \xrightarrow{\text{trop}} & M_g^{\text{trop}} \\
 \downarrow & & \downarrow \\
 A_g & \xrightarrow{\text{trop}} & A_g^{\text{trop}}
 \end{array}$$

This is also proved, with slightly different hypotheses, in [118, Theorem A].

7 Moduli Spaces

7.1 Moduli of Tropical Curves

The moduli space of tropical curves M_g^{trop} has been constructed by numerous authors [1, 33, 64, 87]. In this section, we give a brief description of this object, with an emphasis on applications to classical algebraic geometry.

Given a finite vertex-weighted graph $\mathbf{G} = (G, \omega)$ in the sense of Sect. 4.4, the set of all vertex-weighted metric graphs (Γ, ω) with underlying finite graph G is naturally identified with

$$M_{\mathbf{G}}^{\text{trop}} := \mathbb{R}_{>0}^{|E(G)|} / \text{Aut}(\mathbf{G})$$

with the Euclidean topology. If \mathbf{G}' is obtained from \mathbf{G} by contracting an edge, then we may think of a metric graph in $M_{\mathbf{G}'}^{\text{trop}}$ as a limit of graphs in $M_{\mathbf{G}}^{\text{trop}}$ in which the length of the given edge approaches zero. Similarly, if \mathbf{G}' is obtained from \mathbf{G} by contracting a cycle to a vertex v and augmenting the weight of v by one, we may think of a metric graph in $M_{\mathbf{G}'}^{\text{trop}}$ as a limit of graphs in $M_{\mathbf{G}}^{\text{trop}}$. In this way, we may construct the *moduli space of tropical curves*

$$M_g^{\text{trop}} := \bigsqcup M_{\mathbf{G}}^{\text{trop}},$$

where the union is over all *stable*⁹ vertex-weighted graphs \mathbf{G} of genus g , and the topology is induced by gluing $M_{\mathbf{G}'}^{\text{trop}}$ to the boundary of $M_{\mathbf{G}}^{\text{trop}}$ whenever \mathbf{G}' is a contraction of \mathbf{G} in one of the two senses above.

We note that the moduli space M_g^{trop} is not compact, since edge lengths in a metric graph must be finite and thus there is no limit if we let some edge length tend to infinity. There exists a compactification $\overline{M}_g^{\text{trop}}$, known as the moduli space of *extended tropical curves*, which parameterizes vertex-weighted metric graphs with possibly infinite edges; we refer to [1] for details.

⁹A vertex-weighted finite graph (G, ω) is called *stable* if every vertex of weight zero has valence at least 3.

Let M_g be the (coarse) moduli space of genus g curves and \overline{M}_g its Deligne–Mumford compactification, considered as varieties over \mathbb{C} endowed with the trivial valuation. Points of M_g^{an} can be identified with equivalence classes of points of $M_g(L)$, where L is a complete non-Archimedean field extension of \mathbb{C} (with possibly nontrivial valuation). There is a natural map $\text{Trop} : M_g^{\text{an}} \rightarrow M_g^{\text{trop}}$ which on the level of L -points takes a smooth proper genus g curve C/L to the minimal skeleton of its Berkovich analytification C^{an} . This map extends naturally to a map $\text{Trop} : \overline{M}_g^{\text{an}} \rightarrow \overline{M}_g^{\text{trop}}$.

Let $\Sigma(\overline{M}_g)$ (resp., $\overline{\Sigma}(\overline{M}_g)$) denote the skeleton, in the sense of Thuillier, of M_g^{an} (resp., $\overline{M}_g^{\text{an}}$) with respect to the natural toroidal structure coming from the boundary strata of $\overline{M}_g \setminus M_g$. According to the main result of Abramovich et al. [1], there is a very close connection between the moduli space of tropical curves M_g^{trop} and the Thuillier skeleton $\Sigma(\overline{M}_g)$:

Theorem 7.1 ([1]). *There is a canonical homeomorphism¹⁰*

$$\Phi : \Sigma(\overline{M}_g^{\text{an}}) \rightarrow M_g^{\text{trop}}$$

which extends uniquely to a map

$$\overline{\Phi} : \overline{\Sigma}(\overline{M}_g^{\text{an}}) \rightarrow \overline{M}_g^{\text{trop}}$$

of compactifications in such a way that

$$\begin{array}{ccc} \overline{M}_g^{\text{an}} & \xrightarrow{\overline{P}} & \overline{\Sigma}(\overline{M}_g^{\text{an}}) \\ & \searrow \text{Trop} & \downarrow \overline{\Phi} \\ & & \overline{M}_g^{\text{trop}} \end{array}$$

commutes, where $\overline{P} : \overline{M}_g^{\text{an}} \rightarrow \overline{\Sigma}(\overline{M}_g^{\text{an}})$ is the canonical deformation retraction.

It follows from Theorem 7.1 that the map $\text{Trop} : \overline{M}_g^{\text{an}} \rightarrow \overline{M}_g^{\text{trop}}$ is continuous, proper, and surjective. From this, one easily deduces:

Corollary 7.2. *Let K be a complete and algebraically closed non-Archimedean field with value group \mathbb{R} , and let Γ be a stable metric graph of genus at least 2. Then there exists a curve C over K such that the minimal skeleton of the Berkovich analytic space C^{an} is isometric to Γ .*

¹⁰This homeomorphism is in fact an isomorphism of “generalized cone complexes with integral structure” in the sense of Abramovich et al. [1].

Remark 7.3. A more direct proof of Corollary 7.2, which in fact proves a stronger statement by replacing Γ with an arbitrary metrized complex of curves, and the field K with any complete and algebraically closed non-Archimedean field whose value group contains all edge lengths in some model for Γ , can be found in Theorem 3.24 of [8]. The proof uses formal and rigid geometry. A variant of Corollary 7.2 for discretely valued fields, proved using deformation theory, can be found in Appendix B of [17].

Remark 7.4. Let R be a complete DVR with field of fractions K and infinite residue field κ . The argument in Appendix B of [17] shows that for any finite connected graph G , there exists a regular, proper, flat curve \mathcal{C} over R whose generic fiber is smooth and whose special fiber is a maximally degenerate semistable curve with dual graph G . One should note the assumption here that κ has infinite residue field. In the case where the residue field is finite—for example, when $K = \mathbb{Q}_p$ —the question of which graphs arise in this way remains an open problem. The significance of this problem is its relation to the rational points of the moduli space of curves. For example, the existence of Brill–Noether general curves defined over \mathbb{Q} for large g is a well-known open question. Lang’s conjecture predicts that, for large g , such curves should be contained in a proper closed subset of the moduli space of curves. One suggested candidate for this closed subset is the stable base locus of the canonical bundle, which is known to contain only Brill–Noether special curves.

7.2 Brill–Noether Rank

The motivating problem of Brill–Noether theory is to describe the variety $W_d^r(C)$ parameterizing divisors of a given degree and rank on a curve C . A first step in such a description should be to compute numerical invariants of $W_d^r(C)$, such as its dimension. Our goal is to use the combinatorics of the dual graph Γ to describe $W_d^r(\Gamma)$. Combining this combinatorial description with the Specialization Theorem, we can then hope to understand the Brill–Noether locus of our original curve. One might be tempted to think that the tropical analogue of $\dim W_d^r(C)$ should be $\dim W_d^r(\Gamma)$, but as in the case of linear series, the dimension is not a well-behaved tropical invariant. We note one example of such poor behavior.

Example 7.5. In [93, Theorem 1.1], it is shown that the function that takes a metric graph Γ to $\dim W_d^r(\Gamma)$ is not upper semicontinuous on M_g^{trop} . To see this, the authors construct the following example. Let Γ be the loop of loops of genus 4 depicted in Fig. 11, with edges of length $\ell_1 < \ell_2 < \ell_3$ as pictured. Suppose that $\ell_1 + \ell_2 > \ell_3$. Then $\dim W_3^1(\Gamma) = 1$. If, however, we consider the limiting metric graph Γ_0 as ℓ_1, ℓ_2 , and ℓ_3 approach zero, then on this graph the only divisor of degree 3 and rank 1 is the sum of the three vertices, hence $\dim W_3^1(\Gamma_0) = 0$.




Fig. 11 The metric graph Γ from [93]

The solution to this problem has a very similar flavor to the definition of rank recorded in Definition 4.2. Specifically, given a curve C , consider the incidence correspondence

$$\Phi = \{(p_1, \dots, p_d, D) \in C^d \times W_d^r(C) \mid p_1 + \dots + p_d \in |D|\}.$$

The forgetful map to $W_d^r(C)$ has fibers of dimension r , so $\dim \Phi = r + \dim W_d^r(C)$, and hence the image of Φ in C^d has the same dimension. This suggests the following surrogate for the dimension of $W_d^r(C)$.

Definition 7.6. Let Γ be a metric graph, and suppose that $W_d^r(\Gamma)$ is nonempty. The Brill–Noether rank $w_d^r(\Gamma)$ is the largest integer k such that, for every effective divisor E of degree $r + k$, there exists a divisor $D \in W_d^r(\Gamma)$ such that $|D - E| \neq \emptyset$.

Example 7.7. Note that, in the previous example, although $\dim W_3^1(\Gamma) = 1$, the Brill–Noether rank $w_3^1(\Gamma) = 0$. To see this, it suffices to find a pair of points such that no divisor of degree 3 and rank 1 passes through both points simultaneously. Indeed, it is shown in [93, Theorem 1.9] that no divisor of rank 1 and degree 3 contains $v_3 + w_3$.

The Brill–Noether rank is much better behaved than the dimension of the Brill–Noether locus; for example (cf. [93, Theorem 1.6] and [90, Theorem 5.4]):

Theorem 7.8. *The Brill–Noether rank is upper semicontinuous on the moduli space of tropical curves.*

The Brill–Noether rank also satisfies the following analogue of the Specialization Theorem (cf. [93, Theorem 1.7] and [90, Theorem 5.7]):

Theorem 7.9. *If C is a curve over an algebraically closed field K with nontrivial valuation, and the skeleton of the Berkovich analytic space C^{an} is isometric to Γ , then*

$$\dim W_d^r(C) \leq w_d^r(\Gamma).$$

We note the following generalization of Theorem 4.6.

Corollary 7.10. *Let Γ be a metric graph of genus g . Then $w_d^r(\Gamma) \geq \rho := g - (r + 1) (g - d + r)$.*

Proof. The general theory of determinantal varieties shows that, if $W_d^r(C)$ is nonempty, then its dimension is at least ρ . The result then follows from [84, 86] and Theorem 7.9.

Remark 7.11. It is unknown whether $W_d^r(\Gamma)$ must have local dimension at least ρ . Note, however, that this must hold in a neighborhood of a divisor $D \in \text{Trop } W_d^r(C)$. Hence, if $W_d^r(\Gamma)$ has smaller than the expected local dimension in a neighborhood of some divisor D , then D does not lift to a divisor of rank r on a curve C having Γ as its tropicalization.

8 Metrized Complexes of Curves and Limit Linear Series

In this section we describe the work of Amini and Baker [5] on the Riemann–Roch and Specialization Theorems for divisors on metrized complexes of curves, along with applications to the theory of limit linear series.

8.1 Metrized Complexes of Curves

Metrized complexes of curves can be thought of, loosely, as objects which interpolate between classical and tropical algebraic geometry. More precisely, a *metrized complex of algebraic curves* over an algebraically closed field κ is a finite metric graph Γ together with a fixed model G and a collection of marked complete nonsingular algebraic curves C_v over κ , one for each vertex v of G ; the set \mathcal{A}_v of marked points on C_v is in bijection with the edges of G incident to v . A metrized complex over \mathbb{C} can be visualized as a collection of compact Riemann surfaces connected together via real line segments, as in Fig. 12.

The *geometric realization* $|\mathfrak{C}|$ of a metrized complex of curves \mathfrak{C} is defined as the topological space given by the union of the edges of G and the collection of curves C_v , with each endpoint v of an edge e identified with the corresponding marked point x_v^e (as suggested by Fig. 12). The *genus* of a metrized complex of curves \mathfrak{C} , denoted $g(\mathfrak{C})$, is by definition $g(\mathfrak{C}) = g(\Gamma) + \sum_{v \in V} g_v$, where g_v is the genus of C_v and $g(\Gamma)$ is the first Betti number of Γ .

A *divisor* on a metrized complex of curves \mathcal{C} is an element \mathcal{D} of the free abelian group on $|\mathcal{C}|$. Thus a divisor on \mathcal{C} can be written uniquely as $\mathcal{D} = \sum_{x \in |\mathcal{C}|} a_x x$ where $a_x \in \mathbb{Z}$, all but finitely many of the a_x are zero, and the sum is over all points of $\Gamma \setminus V$ as well as $C_v(\kappa)$ for $v \in V$. The *degree* of \mathcal{D} is defined to be $\sum a_x$.

A *nonzero rational function* f on a metrized complex of curves \mathcal{C} is the data of a rational function $f_\Gamma \in \text{PL}(\Gamma)$ and nonzero rational functions f_v on C_v for each $v \in V$. We call f_Γ the Γ -*part* of f and f_v the C_v -*part* of f . The *divisor* of a nonzero rational function f on \mathcal{C} is defined to be

$$\text{div}(f) := \sum_{x \in |\mathcal{C}|} \text{ord}_x(f) x,$$


where $\text{ord}_x(f)$ is defined as follows¹¹:

- If $x \in \Gamma \setminus V$, then $\text{ord}_x(f) = \text{ord}_x(f_\Gamma)$, where $\text{ord}_x(f_\Gamma)$ is the sum of the slopes of f_Γ in all tangent directions emanating from x .
- If $x \in C_v(\kappa) \setminus \mathcal{A}_v$, then $\text{ord}_x(f) = \text{ord}_x(f_v)$.
- If $x = x_v^e \in \mathcal{A}_v$, then $\text{ord}_x(f) = \text{ord}_x(f_v) + \text{slp}_e(f_\Gamma)$, where $\text{slp}_e(f_\Gamma)$ is the outgoing slope of f_Γ at v in the direction of e .

Divisors of the form $\text{div}(f)$ are called *principal*, and the principal divisors form a subgroup of $\text{Div}^0(\mathcal{C})$, the group of divisors of degree zero on \mathcal{C} . Two divisors in $\text{Div}(\mathcal{C})$ are called *linearly equivalent* if they differ by a principal divisor. Linear equivalence of divisors on \mathcal{C} can be understood rather intuitively in terms of “chip-firing moves” on \mathcal{C} . We refer the reader to Sect. 1.2 of [5] for details.

A divisor $\mathcal{E} = \sum_{x \in |\mathcal{C}|} a_x(x)$ on \mathcal{C} is called *effective* if $a_x \geq 0$ for all x . The *rank* $r_{\mathcal{C}}$ of a divisor $\mathcal{D} \in \text{Div}(\mathcal{C})$ is defined to be the largest integer k such that $\mathcal{D} - \mathcal{E}$ is linearly equivalent to an effective divisor for all effective divisors \mathcal{E} of degree k on \mathcal{C} (so in particular $r_{\mathcal{C}}(\mathcal{D}) \geq 0$ if and only if \mathcal{D} is linearly equivalent to an effective divisor, and otherwise $r_{\mathcal{C}}(\mathcal{D}) = -1$).

Fig. 12 An example of a metrized complex



¹¹Note that our sign convention here for the divisor of a rational function on Γ , which coincides with the one used in [5], is the opposite of the sign convention used in Sect. 3.2, which is also used in a number of other papers in the subject. This should not cause any confusion, but it is good for the reader to be aware of this variability when perusing the literature.

The theory of divisors, linear equivalence, and ranks on metrized complexes of curves generalizes both the classical theory for algebraic curves and the corresponding theory for metric graphs. The former corresponds to the case where G consists of a single vertex v and no edges and $C = C_v$ is an arbitrary smooth curve. The latter corresponds to the case where the curves C_v have genus zero for all $v \in V$. Since any two points on a curve of genus zero are linearly equivalent, it is easy to see that the divisor theories and rank functions on \mathcal{C} and Γ are essentially the same.

The *canonical divisor* on \mathcal{C} is

$$\mathcal{K} = \sum_{v \in V} (K_v + \sum_{w \in \mathcal{A}_v} w),$$

where K_v is a canonical divisor on C_v .

The following result generalizes both the classical Riemann–Roch theorem for algebraic curves and the Riemann–Roch theorem for metric graphs:

Riemann–Roch for Metrized Complexes. *Let \mathcal{C} be a metrized complex of algebraic curves over κ . For any divisor $\mathcal{D} \in \text{Div}(\mathcal{C})$, we have*

$$r_{\mathcal{C}}(\mathcal{D}) - r_{\mathcal{C}}(\mathcal{K} - \mathcal{D}) = \text{deg}(\mathcal{D}) - g(\mathcal{C}) + 1.$$

As with the tropical Riemann–Roch theorem, the proof of this theorem makes use of a suitable notion of *reduced divisors* for metrized complexes of curves. We note that the proof of the Riemann–Roch theorem for metrized complexes uses the Riemann–Roch theorem for algebraic curves and does not furnish a new proof of that result.

8.2 Specialization of Divisors from Curves to Metrized Complexes

Let K be a complete and algebraically closed non-Archimedean field with valuation ring R and residue field κ , and let C be a smooth proper curve over K . As in Sect. 6, there is a metrized complex \mathcal{C} canonically associated with any strongly semistable model \mathcal{C} of C over R . The specialization map Trop defined in Sects. 2 and 3 can be enhanced in a canonical way to a map from divisors on C to divisors on \mathcal{C} . The enhanced specialization map, which by abuse of terminology we continue to denote by Trop, is obtained by linearly extending a map $\tau : C(K) \rightarrow |\mathcal{C}|$. The map τ is defined as follows:

- For $P \in C(K)$ reducing to a smooth point $\text{red}(P)$ of the special fiber C_0 of \mathcal{C} , $\tau(P)$ is just the point $\text{red}(P)$.
- For $P \in C(K)$ reducing to a singular point, $\tau(P)$ is the point $\text{Trop}(P)$ in the relative interior of the corresponding edge of the skeleton Γ of \mathcal{C} .

The motivation for the definitions of $\text{Trop} : C(K) \rightarrow |\mathcal{C}|$ and $\text{div}(f)$ comes in part from the following extension of the Slope Formula (Theorem 6.4):

Proposition 8.1. *Let f be a nonzero rational function on C and let \mathfrak{f} be the corresponding nonzero rational function on \mathfrak{C} , where f_Γ is the restriction to Γ of the piecewise linear function $\log |f|$ on C^{an} and $f_v \in \kappa(C_v)$ for $v \in V$ is the normalized reduction of f to C_v (cf. Sect. 6.1). Then*

$$\text{Trop}(\text{div}(f)) = \text{div}(\mathfrak{f}).$$

In particular, we have $\text{Trop}(\text{Prin}(C)) \subseteq \text{Prin}(\mathfrak{C})$.

The Specialization Theorem from Sect. 4.2 generalizes to metrized complexes as follows:

Specialization Theorem for Metrized Complexes. *For all $D \in \text{Div}(C)$, we have*

$$r_C(D) \leq r_{\mathfrak{C}}(\text{trop}(D)).$$

Since $r_{\mathfrak{C}}(\text{trop}(D)) \leq r_\Gamma(\text{trop}(D))$, the specialization theorem for metrized complexes is a strengthening of the analogous specialization result for metrized graphs. In conjunction with a simple combinatorial argument, this theorem also refines the Specialization Theorem for vertex-weighted graphs.

A simple consequence of the Riemann–Roch and Specialization Theorems for metrized complexes is that for any canonical divisor K_C on C , the divisor $\text{trop}(K_C)$ belongs to the canonical class on \mathfrak{C} . Indeed, the Specialization Theorem shows that $r_{\mathfrak{C}}(\text{trop}(K_C)) \geq g - 1$, while Riemann–Roch shows that a divisor of degree $2g - 2$ and rank at least $g - 1$ must be equivalent to \mathcal{K} .

There is also a version of specialization in which one has *equality* rather than just an inequality. One can naturally associate to a rank r divisor D on C a collection $\mathcal{H} = \{H_v\}_{v \in V}$ of $(r + 1)$ -dimensional subspaces of $\kappa(C_v)$, where H_v is the normalized reduction of $\mathcal{L}(D)$ to C_v (cf. Sect. 6). If $\mathcal{F} = \{F_v\}_{v \in V}$, where F_v is any κ -subspace of the function field $\kappa(C_v)$, then for $\mathcal{D} \in \text{Div}(\mathfrak{C})$ we define the \mathcal{F} -restricted rank of \mathcal{D} , denoted $r_{\mathfrak{C}, \mathcal{F}}(\mathcal{D})$, to be the largest integer k such that for any effective divisor \mathcal{E} of degree k on \mathfrak{C} , there is a rational function \mathfrak{f} on \mathfrak{C} whose C_v -parts f_v belong to F_v for all $v \in V$, and such that $\mathcal{D} - \mathcal{E} + \text{div}(\mathfrak{f}) \geq 0$.

Theorem 8.2 (Specialization Theorem for Restricted Ranks). *With notation as above, the \mathcal{H} -restricted rank of the specialization of D is equal to the rank of D , i.e., $r_{\mathfrak{C}, \mathcal{H}}(\text{trop}(D)) = r$.*

8.3 Connections with the Theory of Limit Linear Series

The theory of linear series on metrized complexes of curves has close connections with the Eisenbud–Harris theory of limit linear series for strongly semistable curves of compact type, and allows one to generalize the basic definitions in the Eisenbud–Harris theory to more general semistable curves. The Eisenbud–Harris theory, which they used to settle a number of longstanding open problems in the theory of

algebraic curves, only applies to a rather restricted class of reducible curves, namely, those of *compact type* (i.e., nodal curves whose dual graph is a tree). It has been an open problem for some time to generalize their theory to more general semistable curves.¹²

Recall that the *vanishing sequence* of a linear series $L = (L, W)$ at $p \in C$, where $W \subset H^0(C, L)$ is the ordered sequence

$$a_0^L(p) < \dots < a_r^L(p)$$

of integers k with the property that there exists some $s \in W$ vanishing to order exactly k at p . For strongly semistable curves of compact type, Eisenbud and Harris define a notion of *crude limit* $\mathfrak{g}_d^r L$ on C_0 , which is the data of a (not necessarily complete) degree d and rank r linear series L_v on C_v for each vertex $v \in V$ with the following property: if two components C_u and C_v of C_0 meet at a node p , then for any $0 \leq i \leq r$,

$$a_i^{L_v}(p) + a_{r-i}^{L_u}(p) \geq d.$$

We can canonically associate to a proper strongly semistable curve C_0 a metrized complex \mathfrak{C} of κ -curves, called the *regularization* of C_0 , by assigning a length of 1 to each edge of G . This is the metrized complex associated with any regular smoothing \mathfrak{C} of C_0 over any discrete valuation ring R with residue field κ .

Theorem 8.3. *Let \mathfrak{C} be the metrized complex of curves associated with a strongly semistable curve C_0/κ of compact type. Then there is a bijective correspondence between the following:*

1. *Crude limit \mathfrak{g}_d^r 's on C_0 in the sense of Eisenbud and Harris.*
2. *Equivalence classes of pairs $(\mathcal{H}, \mathcal{D})$, where $\mathcal{H} = \{H_v\}$, H_v is an $(r + 1)$ -dimensional subspace of $\kappa(C_v)$ for each $v \in V$, and \mathcal{D} is a divisor of degree d supported on the vertices of \mathfrak{C} with $r_{\mathfrak{C}, \mathcal{H}}(\mathcal{D}) = r$. Here we say that $(\mathcal{H}, \mathcal{D}) \sim (\mathcal{H}', \mathcal{D}')$ if there is a rational function f on \mathfrak{C} such that $\mathcal{D}' = \mathcal{D} + \text{div}(f)$ and $H_v = H'_v \cdot f_v$ for all $v \in V$, where f_v denotes the C_v -part of f .*

Theorem 8.3, combined with the Riemann–Roch theorem for metrized complexes of curves, provides a new proof of the fact, originally established in [58], that limit linear series satisfy analogues of the classical theorems of Riemann and Clifford. The point is that $r_{\mathfrak{C}, \mathcal{H}}(\mathcal{D}) \leq r_{\mathfrak{C}}(\mathcal{D})$ for all $\mathcal{D} \in \text{Div}(\mathfrak{C})$, and therefore upper bounds on $r_{\mathfrak{C}}(\mathcal{D})$ which follow from Riemann–Roch imply corresponding upper bounds on the restricted rank $r_{\mathfrak{C}, \mathcal{H}}(\mathcal{D})$.

Motivated by Theorem 8.3, Amini and Baker propose the following definition.

Definition 8.4. Let C_0 be a strongly semistable (but not necessarily compact type) curve over κ with regularization \mathfrak{C} . A *limit \mathfrak{g}_d^r* on C_0 is an equivalence class of pairs

¹²Brian Osserman [107] has recently proposed a different framework for doing this.

$(\{H_v\}, \mathcal{D})$ as above, where H_v is an $(r + 1)$ -dimensional subspace of $\kappa(C_v)$ for each $v \in V$, and \mathcal{D} is a degree d divisor on \mathcal{C} with $r_{\mathcal{C}, \mathcal{H}}(\mathcal{D}) = r$.

As partial additional justification for Definition 8.4, Amini and Baker prove, using specialization, that a \mathfrak{g}_d^r on the smooth general fiber C of a semistable family \mathcal{C} gives rise in a natural way to a crude limit \mathfrak{g}_d^r on the central fiber.

Part 3: Applications

In this part, we discuss several recent applications of tropical Brill–Noether theory to problems in algebraic and arithmetic geometry. These sections are largely independent of each other, so the reader should be able to peruse them according to his or her interest.

9 Applications of Tropical Linear Series to Classical Brill–Noether Theory

Recent years have witnessed several applications of tropical Brill–Noether theory to problems in classical algebraic geometry. In this section, we survey the major recent developments in the field.

9.1 The Brill–Noether Theorem

The Brill–Noether theorem predicts the dimension of the space $W_d^r(C)$ parameterizing divisor classes (or, equivalently, complete linear series) of a given degree and rank on a general curve C .

Brill–Noether Theorem ([66]). *Let C be a general curve of genus g over \mathbb{C} . Then $W_d^r(C)$ has pure dimension $\rho(g, r, d) = g - (r + 1)(g - d + r)$, if this is nonnegative, and is empty otherwise.*

The original proof of the Brill–Noether theorem, due to Griffiths and Harris, involves a subtle degeneration argument [66]. The later development of limit linear series by Eisenbud and Harris led to a simpler proof of this theorem [57, 58]. The literature contains several other proofs, some of which work in any characteristic. One that is often referenced is due to Lazarsfeld, because rather than using degenerations, Lazarsfeld’s argument involves vector bundles on K3 surfaces [89].

The first significant application of tropical Brill–Noether theory was the new proof of the Brill–Noether theorem by Cools, Draisma, Payne, and Robeva [48], which successfully realized the program laid out in [17]. In [48], the authors




Fig. 13 The graph Γ

consider the family of graphs pictured in Fig. 13, colloquially known as the chain of loops.¹³ The edge lengths are further assumed to be generic, which in this case means that, if ℓ_i, m_i are the lengths of the bottom and top edges of the i th loop, then ℓ_i/m_i is not equal to the ratio of two positive integers whose sum is less than or equal to $2g - 2$.

Using Theorem 4.6 as the only input from algebraic geometry, the authors of [48] employ an intricate combinatorial argument to prove the following:

Theorem 9.1 ([48]). *Let C be a smooth projective curve of genus g over a discretely valued field with a regular, strongly semistable model whose special fiber is a generic chain of loops Γ . Then $\dim W_d^r(C) = \rho(g, r, d)$ if this number is nonnegative, and $W_d^r(C) = \emptyset$ otherwise.*

We note that such a curve C exists by Corollary 7.2. The Brill–Noether theorem (over an arbitrary algebraically closed field) then follows from Theorem 9.1 using the theory of Brill–Noether rank discussed in Sect. 7.2.

In fact, [48] proves more. Theorem 4.6 of [48] completely describes $W_d^r(\Gamma)$, explicitly classifying all divisors of given degree and rank on a generic chain of loops. Indeed, it is shown that $W_d^r(\Gamma)$ is a union of ρ -dimensional tori. The set of tori is in bijection with the so-called lingering lattice paths, which in turn are in bijection with standard Young tableaux on a rectangle with $r + 1$ columns and $g - d + r$ rows containing the numbers $1, \dots, g$. From this, one can compute the number of tori to be

$$\binom{g}{\rho} (g - \rho)! \prod_{i=0}^r \frac{i!}{(g - d + r + i)!}$$

if $\rho \geq 0$, and 0 if $\rho < 0$.

We briefly discuss the argument here. Given an effective divisor D , we may assume that D is v_1 -reduced. The divisor D then has some number d_1 of chips at

¹³In fact, they consider the graph in which the lengths of the bridge edges between the loops are all zero. There is, however, a natural rank-preserving isomorphism between the Jacobian of a metric graph with a bridge and the Jacobian of the graph in which that bridge has been contracted, so their argument works equally well in this case. We consider the graph with bridges because of its use in [78, 79].




Fig. 14 A decomposition of Γ

v_1 , and by Dhar’s burning algorithm D has at most 1 chip on each of the half-open loops γ_k pictured in Fig. 14, and no chips on the half-open bridges br_k .

The associated *lingering lattice path* is a sequence of vectors $p_i \in \mathbb{Z}^r$, starting at $p_1 = (d_1, d_1 - 1, \dots, d_1 - r + 1)$, with the i th step given by

$$p_{i+1} - p_i = \left\{ \begin{array}{ll} (-1, -1, \dots, -1) & \text{if } D \text{ has no chip on } \gamma_i \\ e_j & \text{if } D \text{ has a chip on } \gamma_i, \text{ the distance from } v_i \\ & \text{to this chip is precisely } (p_i(j) + 1)m_{i+1}, \\ & \text{and both } p_i \text{ and } p_i + e_j \text{ are in } \mathcal{W} \\ 0 & \text{otherwise} \end{array} \right\}$$

Here, the distance from v_i is in the counterclockwise direction. Since the chip lies on a circle of circumference $m_i + \ell_i$, this distance should be understood to be modulo $(m_i + \ell_i)$. The symbols e_0, \dots, e_{r-1} represent the standard basis vectors in \mathbb{Z}^r and \mathcal{W} is the open Weyl chamber

$$\mathcal{W} = \{y \in \mathbb{Z}^r | y_0 > y_1 > \dots > y_{r-1} > 0\}.$$

The steps where $p_{i+1} = p_i$ are known as *lingering steps*. The basic idea of the lingering lattice path is as follows. By Theorem 5.7, the set $\{v_1, v_2, \dots, v_g, w_g\}$ is rank-determining. Hence, if D fails to have rank r , there is an effective divisor E of degree r , supported on these vertices, such that $|D - E| = \emptyset$. Starting with the v_1 -reduced divisor D , we move chips to the right and record the v_i -degree of the equivalent v_i -reduced divisor. The number $p_i(j)$ is then the minimum, over all effective divisors E of degree j supported at v_1, \dots, v_i , of the v_i -degree of the v_i -reduced divisor equivalent to $D - E$. From this it follows that, when D has rank at least r , we must have $p_i(j) \geq r - j$, so the corresponding lingering lattice path must lie in the open Weyl chamber \mathcal{W} .

The corresponding tableau is constructed by placing the moves in the direction e_j in the i th column of the rectangle, and the moves in the direction $(-1, \dots, -1)$ in the last column. If the k th step is lingering, then the integer k does not appear in the tableau. Given this description, we see that each tableau determines the existence and position of the chip on the half-open loop γ_k if and only if the integer k appears in the tableau. Otherwise, the chip on the k th loop is allowed to move freely. The number of chips that are allowed to move freely is therefore $\rho = g - (r + 1)(g - d + r)$. Indeed, we see that not only is the Brill–Noether rank $w'_d(\Gamma)$ equal to ρ , but in fact

$\dim W_d^r(\Gamma) = \rho$ as well. Theorem 9.1 then follows from the specialization result for Brill–Noether rank, Theorem 7.9.

9.2 The Gieseker–Petri Theorem

Assume that $\rho(g, r, d) \geq 0$. The variety $W_d^r(C)$ is singular along $W_d^{r+1}(C)$. Blowing up along this subvariety yields the variety $\mathcal{G}_d^r(C)$ parameterizing (not necessarily complete) linear series of degree d and rank r on C . A natural generalization of the Brill–Noether theorem is the following:

Gieseker–Petri Theorem ([65]). *Let C be a general curve of genus g . If $\rho(g, r, d) \geq 0$, then $\mathcal{G}_d^r(C)$ is smooth of dimension $\rho(g, r, d)$.*

It is a standard result, following [12, Sect. IV.4], that the Zariski cotangent space to $\mathcal{G}_d^r(C)$ at a point corresponding to a complete linear series $\mathcal{L}(D)$ is naturally isomorphic to the cokernel of the adjoint multiplication map

$$\mu_D : \mathcal{L}(D) \otimes \mathcal{L}(K_C - D) \rightarrow \mathcal{L}(K_C).$$

Thus the cotangent space has dimension $\rho(g, r, d) + \dim \ker \mu_D$, and in particular, $\mathcal{G}_d^r(C)$ is smooth of dimension $\rho(g, r, d)$ at such a point if and only if the multiplication map μ_D is injective. More generally, if $P \in \mathcal{G}_d^r(C)$ corresponds to a possibly incomplete linear series $W \subset \mathcal{L}(D)$, then $\mathcal{G}_d^r(C)$ is smooth of dimension $\rho(g, r, d)$ at P if and only if the multiplication map $W \otimes \mathcal{L}(K_C - D) \rightarrow \mathcal{L}(K_C)$ is injective. One deduces that the Gieseker–Petri theorem is equivalent to the assertion that if C is a general curve of genus g , then μ_D is injective for all divisors D on C .

A recent application of tropical Brill–Noether theory is the following result [78], which yields a new proof of the Gieseker–Petri theorem:

Theorem 9.2 ([78]). *Let C be a smooth projective curve of genus g over a discretely valued field with a regular, strongly semistable model whose special fiber is a generic chain of loops Γ . Then the multiplication map*

$$\mu_D : \mathcal{L}(D) \otimes \mathcal{L}(K_C - D) \rightarrow \mathcal{L}(K_C)$$

is injective for all divisors D on C .

The argument has much in common with the tropical proof of the Brill–Noether theorem, using the same metric graph with the same genericity conditions on edge lengths. The new ingredient is the idea of tropical independence, as defined in Sect. 5.3. Given a divisor $D \in W_d^r(C)$, the goal is to find functions

$$f_0, \dots, f_r \in \text{trop}(\mathcal{L}(D))$$

$$g_0, \dots, g_{g-d+r-1} \in \text{trop}(\mathcal{L}(K_C - D))$$

such that $\{f_i + g_j\}_{i,j}$ is tropically independent.

There is a dense open subset of $W_d^r(\Gamma)$ consisting of divisors D with the following property: given an integer $0 \leq i \leq r$, there exists a unique divisor $D_i \sim D$ such that

$$D_i - iw_g - (r - i)v_1 \geq 0.$$

These are the divisors referred to as *vertex-avoiding* in [41].

We first describe the proof of Theorem 9.2 in the case that D is vertex-avoiding. If D is the specialization of a divisor $\mathcal{D} \in W_d^r(C)$, and $p_1, p_g \in C$ are points specializing to v_1, w_g , respectively, then there exists a divisor $\mathcal{D}_i \sim \mathcal{D}$ such that

$$\mathcal{D}_i - ip_g - (r - i)p_1 \geq 0,$$

and, by the uniqueness of D_i, \mathcal{D}_i must specialize to D_i . It follows that there is a function $f_i \in \text{trop}(\mathcal{L}(\mathcal{D}))$ such that $\text{div}(f_i) = D_i - D$.

For this open subset of divisors, the argument then proceeds as follows. By the classification in [48], the divisor D_i fails to have a chip on the k th loop if and only if the integer k appears in the i th column of the corresponding tableau. The adjoint divisor $E = K_\Gamma - D$ corresponds to the transpose tableau [2, Theorem 39], so the divisor $D_i + E_j$ fails to have a chip on the k th loop if and only if k appears in the (i, j) position of the tableau. Since for each k at most one of these divisors fails to have a chip on the k th loop, we see that if

$$\theta = \min\{f_i + g_j + b_{i,j}\}$$

occurs at least twice at every point of Γ , then the divisor

$$\Delta = \text{div}(\theta) + K_\Gamma$$

must have a chip on the k th loop for all k .

To see that this is impossible, let p_k be a point of Δ in γ_k , and let

$$D' = p_1 + \cdots + p_g.$$

Then by construction $K_\Gamma - D'$ is equivalent to an effective divisor, so by the tropical Riemann–Roch theorem we see that $r(D') \geq 1$. On the other hand, Dhar’s burning algorithm shows that D' is universally reduced, so by Proposition 5.16 we have $r(D') = 0$, a contradiction.

It is interesting to note that this obstruction is, at heart, combinatorial. Unlike the earlier proofs via limit linear series, which arrive at a contradiction by constructing a canonical divisor of impossible *degree* (larger than $2g - 2$), this argument arrives at a contradiction by constructing a canonical divisor of impossible *shape*.

The major obstacle to extending this argument to the case where D is not vertex-avoiding is that the containment $\text{trop}(\mathcal{L}(D)) \subseteq R(\text{Trop}(D))$ is often strict. Given an arbitrary divisor $D \in W_d^r(C)$ and function $f \in R(\text{Trop}(D))$, it is difficult to determine whether f is the specialization of a function in $\mathcal{L}(D)$. To avoid this issue, the authors

make use of a patching construction, gluing together tropicalizations of different rational functions in a fixed algebraic linear series on different parts of the graph, to arrive at a piecewise linear function in $R(K_\Gamma)$ that may not be in $\text{trop}(\mathcal{L}(K_C))$. Once this piecewise linear function is constructed, the argument proceeds very similarly to the vertex-avoiding case.

9.3 The Maximal Rank Conjecture

One of the most well-known open problems in Brill–Noether theory is the Maximal Rank Conjecture, which predicts the Hilbert function for sufficiently general embeddings of sufficiently general curves. This conjecture is attributed to Noether in [11, p. 4] (see [105, Sect. 8] and [43, pp. 172–173] for details), was studied classically by Severi [112, Sect. 10], and popularized by Harris [71, p. 79].

Maximal Rank Conjecture. *Fix nonnegative integers g, r, d , let C be a general curve of genus g , and let $V \subset \mathcal{L}(D)$ be a general linear series of rank r and degree d on C . Then the multiplication maps*

$$\mu_m : \text{Sym}^m V \rightarrow \mathcal{L}(mD)$$

have maximal rank for all m . That is, each μ_m is either injective or surjective.

While the Maximal Rank Conjecture remains open in general, several important cases are known [27, 61, 115, 119]. For example, it is shown in [27] that the Maximal Rank Conjecture holds in the non-special range $d \geq g + r$. When $d < g + r$, the general linear series of degree d and rank r on a general curve is complete, and for this reason, most of the work in the subject focuses on the case where $V = \mathcal{L}(D)$. We note that the arguments of Ballico and Ellia [27] and Farkas [61] involve degenerations to unions of two curves that meet in more than one point. Since such curves are not of compact type, the arguments do not make use of limit linear series.

In [79], tropical Brill–Noether theory is used to prove the $m = 2$ case of the Maximal Rank Conjecture.

Theorem 9.3 ([79]). *Let C be a smooth projective curve of genus g over a discretely valued field with a regular, strongly semistable model whose special fiber is a generic chain of loops Γ . For a given r and d , let D be a general divisor of rank r and degree d on C . Then the multiplication map*

$$\mu_2 : \text{Sym}^2 \mathcal{L}(D) \rightarrow \mathcal{L}(2D)$$

has maximal rank.

The genericity conditions placed on the edge lengths of Γ in Theorem 9.3 are stricter than those appearing in the tropical proofs of the Brill–Noether and

Gieseker–Petri theorems. First, the bridges between the loops are assumed to be much longer than the loops themselves, and second, one must assume that certain integer linear combinations of the edge lengths do not vanish.

A simplifying aspect of the Maximal Rank Conjecture is that it concerns a general, rather than arbitrary, divisor. It therefore suffices to prove that the maximal rank condition holds for a single divisor of the given degree and rank on C . The main result of Cartwright et al. [41] is that every divisor on the generic chain of loops is the specialization of a divisor of the same rank on C . We are therefore free to choose whatever divisor we wish to work with, and in particular we may choose one of the vertex-avoiding divisors described in the previous section. Recall that, if $D \in W_d^r(\Gamma)$ is vertex-avoiding, then we have an explicit set of piecewise linear functions $f_i \in R(D)$ that are tropicalizations of a basis for the linear series on the curve C . The goal, in the case where the multiplication map is supposed to be injective, is to show that the set $\{f_i + f_j\}_{i \leq j}$ is tropically independent. In the surjective case, we must choose a subset of the appropriate size, and then show that this subset is tropically independent.

The basic idea of the argument is as follows. Assume that

$$\theta = \min\{f_i + f_j + b_{i,j}\}$$

occurs at least twice at every point of Γ , and consider the divisor

$$\Delta = \text{div}(\theta) + 2D.$$

To arrive at a contradiction, one studies the degree distribution of the divisor Δ across the loops of Γ . More precisely, one defines

$$\delta_k := \text{deg}(\Delta|_{\gamma_k}).$$

The first step is to show that $\delta_k \geq 2$ for all k . One then identifies intervals $[a, b]$ for which this inequality must be strict for at least one $k \in [a, b]$. As one proceeds from left to right across the graph, one encounters such intervals sufficiently many times to obtain $\text{deg } \Delta > 2 \text{ deg } D$, a contradiction.

10 Lifting Problems for Divisors on Metric Graphs

In this section we discuss the lifting problem in tropical Brill–Noether theory: given a divisor of rank r on a metric graph Γ , when is it the tropicalization of a rank r divisor on a smooth curve C ? There are essentially two formulations of this problem, one in which the curve C is fixed, and one in which it is not.

Throughout this section, we let K be a complete and algebraically closed nontrivially valued non-Archimedean field.

Question 10.1. Given a metric graph Γ and a divisor D on Γ , under what conditions do there exist a curve C/K (together with a semistable model R) and a divisor of the same rank as D tropicalizing to Γ and D , respectively?

Question 10.2. Given a curve C/K (together with a semistable model R) tropicalizing to a metric graph Γ , and given a divisor D on Γ , under what conditions does there exist a divisor on C of the same rank as D tropicalizing to D ?

These are very difficult questions. Even for the earlier theory of limit linear series on curves of compact type, the analogous questions remain open. A partial answer in that setting is given by the Regeneration Theorem of Eisenbud and Harris [58], which says that if the space of limit linear series has local dimension equal to the Brill–Noether number ρ , then the given limit linear series lifts in any one-parameter smoothing. At the time of writing, there is no corresponding theorem in the tropical setting.¹⁴

10.1 Specialization of Hyperelliptic Curves

One of the first results concerning lifting of divisors is the classification of vertex-weighted metric graphs that are the specialization of a hyperelliptic curve. Recall from Sect. 4.4 that given a curve C/K and a semistable model \mathcal{C}/R for C , there is a natural way to associate to \mathcal{C} a vertex-weighted metric graph (Γ, ω) . We call such a pair *minimal* if there is no vertex v with $\text{val}(v) = 1$ and $\omega(v) = 0$.

Theorem 10.3 ([9, 34]). *Let (Γ, ω) be a minimal vertex-weighted metric graph. There is a smooth projective hyperelliptic curve over a discretely valued field with a regular, strongly semistable model whose special fiber has dual graph Γ if and only if the following conditions hold:*


- (HYP1) *there exists an involution s on Γ such that the quotient Γ/s is a tree and $s(v) = v$ for all $v \in \Gamma$ with $\omega(v) > 0$ and*
- (HYP2) *for every point $v \in \Gamma$, the number of bridge edges adjacent to v is at most $2\omega(v) + 2$.*

Kawaguchi and Yamaki show moreover that, when Γ satisfies these conditions, there is a smoothing C for which every divisor on Γ lifts to a divisor of the same rank on C [83].

We outline the necessity of the conditions above in the special case where $\omega = 0$, which is equivalent to requiring that $g(C) = g(\Gamma)$. Note that if C is a hyperelliptic curve, then by the Specialization Theorem any divisor of degree 2 and rank 1 on C specializes to a divisor D of rank at least 1 on Γ , and by tropical Clifford’s theorem D must have rank exactly 1. Now, if $P \in \Gamma$ is not contained in a bridge, then $|P| = \{P\}$.

¹⁴Amini has apparently made substantial progress in this direction.

Fig. 15 A hyperelliptic metric graph of genus 3 that is not the skeleton of any hyperelliptic curve of genus 3



On the other hand, if $P \in \Gamma$ is contained in a bridge, a simple analysis reveals that $D \sim 2P$. In this way we obtain an involution s on Γ mapping each point P to the P -reduced divisor equivalent to $D - P$.

To see why (HYP2) holds, note that for each type-2 point $v \in \Gamma$, the linear series of degree 2 and rank 1 on C specializes to a linear series of degree 2 and rank 1 on the corresponding curve C_v . Each of the bridges adjacent to v corresponds to ramification points of this linear series, but such a linear series has only $2g(C_v) + 2 = 2$ ramification points.

To see that the conditions (HYP1) and (HYP2) are sufficient requires significantly more work.

Example 10.4. Consider the metric graph Γ pictured in Fig. 15, consisting of a tree with a loop attached to each leaf, with all edge lengths being arbitrary. Then Γ is hyperelliptic, because any divisor of degree 2 supported on the tree has rank 1. On the other hand, this graph is not the dual graph of the limit of any family of genus 3 hyperelliptic curves, because the vertex of valence 3 in the tree is adjacent to more than 2 bridges.

For metric graphs of higher gonality, the lifting problem is significantly harder. In [97], Luo and Manjunath describe an algorithm for smoothability of rank one generalized limit linear series on metrized complexes.

10.2 Lifting Divisors on the Chain of Loops

For some specific families of graphs, such as the chain of loops discussed Sect. 9, one can show that the lifting problem is unobstructed.

Theorem 10.5 ([41]). *Let C/K be a smooth projective curve of genus g . If the dual graph of the central fiber of some regular model of C is isometric to a generic chain of loops Γ of genus g , then every divisor class on Γ that is rational over the value group of K lifts to a divisor class of the same rank on C .*

The general strategy for proving Theorem 10.5 is to study the Brill–Noether loci as subschemes

$$W_d^r(C) \subset \text{Jac}(C).$$

Since C is maximally degenerate, the universal cover of $\text{Jac}(C)^{\text{an}}$ gives a uniformization

$$T^{\text{an}} \rightarrow \text{Jac}(C)^{\text{an}}$$

by an algebraic torus T of dimension g . The tropicalization of this torus is the universal cover of the skeleton of $\text{Jac}(C)$, which as discussed in Sect. 6 is canonically identified with the tropical Jacobian of Γ [20].

A key tool in the proof of Theorem 10.5 is Rabinoff’s lifting theorem [110], which can be applied to the analytic preimages in T of algebraic subschemes of $\text{Jac}(C)$. This lifting theorem says that isolated points in complete intersections of tropicalizations of analytic hypersurfaces lift to points in the analytic intersection with appropriate multiplicities. This theorem can be applied to translates of the preimage of the theta divisor $\Theta_\Gamma = W_{g-1}^0(\Gamma)$, as follows.

When Γ is the generic chain of loops, one can use the explicit description of $W_d^r(\Gamma)$ from [48] to produce explicit translates of Θ_C whose tropicalizations intersect transversally and locally cut out $W_d^r(\Gamma)$. By intersecting with ρ additional translates of Θ_C , one obtains an isolated point in a tropical complete intersection, to which we may apply Rabinoff’s lifting theorem. This complete intersection is typically larger than $W_d^r(\Gamma)$, but the argument shows that the tropicalization map from a 0-dimensional slice of $W_d^r(C)$ to the corresponding slice of $W_d^r(\Gamma)$ is injective. Using again the explicit description of $W_d^r(\Gamma)$, one then shows that the two finite sets have the same cardinality, and hence the map is bijective.

Remark 10.6. As mentioned in the section on the Maximal Rank Conjecture, Theorem 10.5 is one of the key ingredients in the proof of Theorem 9.3 (the Maximal Rank Conjecture for quadrics). In particular, in order to show that the maximal rank condition holds for a generic line bundle of a given degree and rank, it suffices to show that it holds for a single line bundle. Since every divisor of a given rank on the chain of loops lifts to a line bundle on C of the same rank, one is free to work with any divisor of this rank on the chain of loops.

10.3 Examples of Divisors That Do Not Lift

Among the results on lifting divisors, there is a plethora of examples of divisors that do not lift. For example, in [53], Coppens defines a base-point free divisor on a metric graph Γ to be a divisor D such that $r(D - p) < r(D)$ for all $p \in \Gamma$. He then shows that the Clifford and Riemann–Roch bounds are the only obstructions to the existence of base-point free divisors on metric graphs of arbitrary genus. This is in contrast to the case of algebraic curves, where, for example, a curve of genus greater than 6 cannot have a base-point free divisor of degree 5 and rank 2.

Another example of divisors that do not lift comes from the theory of matroids.

Theorem 10.7 ([40]). *Let M be any rank 3 matroid. Then there exist a graph G_M and a rank 2 divisor D_M on G_M such that, for any infinite field k , there are a curve C over $k((t))$ (together with a semistable model \mathcal{C} of C over $k[[t]]$) and a rank 2 divisor on C tropicalizing to G_M and D_M , respectively, if and only if M is realizable over k .*

Combining this with the scheme-theoretic analogue of Mnëv universality, due to Lafforgue [88], one obtains the following.

Corollary 10.8 ([40]). *Let X be a scheme of finite type over $\text{Spec } \mathbb{Z}$. Then there exist a graph G and a rank 2 divisor D on G such that, for any infinite field k , G , and D are the tropicalizations of a curve $C/k((t))$ and a rank 2 divisor on C if and only if X has a k -point.*

In other words, the obstructions to lifting over a valued field of the form $k((t))$ are essentially as general as possible.

Cartwright’s construction is as follows. Recall that a rank 3 simple matroid on a finite set E consists of a collection of subsets of E , called *flats*, such that every pair of elements is contained in exactly one flat. (Here we are abusing language, using the word flat to refer to the maximal, or rank 2, flats.) The bipartite graph G_M is the *Levi graph* of the matroid M , where the vertices correspond to elements and flats, and there is an edge between two vertices if the corresponding element is contained in the corresponding flat. The divisor D_M is simply the sum of the vertices corresponding to elements of E . A combinatorial argument then shows that the rank of D_M is precisely 2.

If M is realizable over k , then by definition, there exists a configuration of lines in \mathbb{P}_k^2 where the lines correspond to the elements of E , and the flats correspond to points where two or more of the lines intersect. If we blow up the plane at the intersection points, the dual graph of the resulting configuration is the Levi graph G_M , and the pullback of the hyperplane class specializes to the divisor D_M . After some technical deformation arguments, one then sees that the pair (G_M, D_M) admits a lifting when M is realizable over k .

For the converse, one must essentially show that the above construction is the only possibility. That is, if \mathcal{C} is a regular semistable curve over $k[[t]]$, the dual graph of the central fiber is G_M and the divisor D_M is the specialization of a rank 2 divisor on \mathcal{C} , then in fact the image of the central fiber under the corresponding linear series must provide a realization of the matroid M in \mathbb{P}_k^2 .

11 Bounding the Number of Rational Points on Curves

By Faltings’ theorem (née the Mordell conjecture), if C is a curve of genus $g \geq 2$ over a number field K , then the set $C(K)$ of rational points on C is finite. Shortly after Faltings proved this theorem, Vojta published a new proof which furnishes an effective upper bound on the number of points in $C(K)$. However, the Vojta bound is completely theoretical—to our knowledge no one has ever written down the bound

explicitly (and the bound is surely quite far from optimal). None of the existing proofs of the Mordell conjecture gives an algorithm—even in theory!—to compute the set $C(K)$. And in practice the situation is even worse—it seems safe to say that no one has ever used the Faltings or Vojta proofs of the Mordell conjecture to compute $C(K)$ in a single nontrivial example.

11.1 *The Katz–Zureick-Brown Refinement of Coleman’s Bound*

One of the first significant results in the direction of the Mordell conjecture was Chabauty’s theorem that $C(K)$ is finite provided that the rank of the finitely generated abelian group $\text{Jac}(C)(K)$ is less than g . Much later, Coleman used his theory of p -adic integration to give an effective upper bound on $C(K)$ in this situation. Coleman’s bound has the advantage of being sharp in certain cases, and the method of proof can be used to compute $C(K)$ in a wide range of concrete examples. For simplicity, we state the results in this section for $K = \mathbb{Q}$ only, but everything extends with minor modifications to curves over a number field K . Coleman’s theorem is as follows.

Theorem 11.1 ([50]). *Let C be a curve of genus g over \mathbb{Q} , and suppose that the Mordell–Weil rank r of $\text{Jac}(C)(\mathbb{Q})$ is strictly less than the genus g . Then for every prime $p > 2g$ of good reduction for C , we have*

$$\#C(\mathbb{Q}) \leq \#C(\mathbb{F}_p) + 2g - 2. \quad (2)$$

Coleman’s theorem was subsequently strengthened in different ways. In [95], Lorenzini and Tucker (see also McCallum–Poonen [99]) generalized Theorem 11.1 to primes of bad reduction, replacing $C(\mathbb{F}_p)$ in (2) by the smooth \mathbb{F}_p -points of the special fiber of the minimal proper regular model for C over \mathbb{Z}_p . Stoll replaced the quantity $2g - 2$ in (2) by $2r$ when C has good reduction at p , and asked if this improvement could be established in the bad reduction case as well. Stoll’s question was answered affirmatively by Katz and Zureick-Brown in [80] by supplementing Stoll’s method with results from the theory of linear series on tropical curves:

Theorem 11.2 ([80]). *Let C be a curve of genus g over \mathbb{Q} and suppose that the rank r of $\text{Jac}(C)(\mathbb{Q})$ is less than g . Then for every prime $p > 2r + 2$, we have*

$$\#C(\mathbb{Q}) \leq \#C^{\text{sm}}(\mathbb{F}_p) + 2r,$$

where C denotes the minimal proper regular model of C over \mathbb{Z}_p .

In order to explain the relevance of linear series on tropical curves to such a result, we need to briefly explain the basic ideas underlying the previous work of Coleman et al. Let us first outline a proof of Theorem 11.1. Fix a rational point $P \in C(\mathbb{Q})$

(if no such point exists, then the theorem is vacuously true) and let $\iota : C \hookrightarrow J$ be the corresponding Abel–Jacobi embedding. Coleman’s theory of p -adic integration of 1-forms associates to each $\omega \in H^0(C, \Omega_C^1)$ and $Q \in C(\mathbb{Q}_p)$ a (definite) p -adic integral $\int_p^Q \omega \in \mathbb{Q}_p$, obtained by pulling back a corresponding p -adic integral on J via the map ι . Locally on C , such p -adic integrals can be computed by formally integrating a power series expansion $f_\omega(T)$ for ω with respect to a local parameter T on some residue disc U . One can show fairly easily that the p -adic closure $\overline{J(\mathbb{Q})}$ of $J(\mathbb{Q})$ in $J(\mathbb{Q}_p)$ has dimension at most r as a p -adic manifold. The formalism of Coleman’s theory implies that forcing the p -adic integral of a 1-form on J to vanish identically on $\overline{J(\mathbb{Q})}$ imposes at most r linear conditions on $H^0(J, \Omega_J^1)$. The functoriality of Coleman integration implies that the \mathbb{Q}_p -vector space V_{chab} of all $\omega \in H^0(C, \Omega_C^1)$ such that $\int_p^Q \omega = 0$ for all $Q \in C(\mathbb{Q})$ has dimension at least $g - r > 0$.

The condition $p > 2g$ implies, by a p -adic analogue of Rolle’s theorem which can be proved in an elementary way with Newton polygons, that if $f_\omega(T)$ has n zeroes on U then $\int f_\omega(T) dT$ has at most $n + 1$ zeroes on U . Using this observation, Coleman deduces, by summing over all residue classes, that if ω is a nonzero 1-form in $H^0(C, \Omega_C^1)$ vanishing on all of $C(\mathbb{Q})$ then

$$\#C(\mathbb{Q}) \leq \sum_{\overline{Q} \in \overline{C}(\mathbb{F}_p)} (1 + \text{ord}_{\overline{Q}} \overline{\omega}),$$

where $\overline{\omega}$ denotes the reduction of ω to \overline{C} . Since the 1-form $\overline{\omega}$ on \overline{C} has a total of $2g - 2$ zeros counting multiplicity, we have

$$\sum_{\overline{Q} \in \overline{C}(\mathbb{F}_p)} \text{ord}_{\overline{Q}} \overline{\omega} \leq 2g - 2,$$

which yields Coleman’s bound.

Stoll observed in [113] that one could do better than this by adapting the differential ω to the point \overline{Q} rather than using the same differential ω for all residue classes. Define the *Chabauty divisor*

$$D_{\text{chab}} = \sum_{\overline{Q} \in \overline{C}(\mathbb{F}_p)} n_{\overline{Q}}(\overline{Q}),$$

where $n_{\overline{Q}}$ is the minimum over all nonzero ω in V_{chab} of $\text{ord}_{\overline{Q}} \overline{\omega}$, and let d be the degree of D_{chab} . Since D_{chab} and $K_{\overline{C}} - D_{\text{chab}}$ are both equivalent to effective divisors, Clifford’s inequality (applied to the smooth proper curve \overline{C}) implies that

$$h^0(K_{\overline{C}} - D_{\text{chab}}) - 1 \leq \frac{1}{2}(2g - 2 - d).$$

On the other hand, the semicontinuity of $h^0(D) = r(D) + 1$ under specialization shows that $h^0(D_{\text{chab}}) \geq \dim V_{\text{chab}} \geq g - r$. Combining these inequalities gives

$$g - r - 1 \leq \frac{1}{2}(2g - 2 - d)$$

and thus $d \leq 2r$, giving Stoll’s refinement of Coleman’s bound.

Lorenzini and Tucker [95] had shown earlier that one can generalize Coleman’s bound to the case of bad reduction as follows. Since points of $C(\mathbb{Q})$ specialize to the set $\bar{C}^{\text{sm}}(\mathbb{F}_p)$ of smooth \mathbb{F}_p -points on the special fiber of C under the reduction map, one obtains by an argument similar to the one above the bound

$$\#C(\mathbb{Q}) \leq \sum_{\bar{Q} \in \bar{C}^{\text{sm}}(\mathbb{F}_p)} (1 + n_{\bar{Q}}), \tag{3}$$

where $\bar{\omega}$ denotes the reduction of ω to the unique irreducible component of the special fiber of C containing \bar{Q} . Choosing a nonzero $\omega \in V_{\text{chab}}$ as in Coleman’s bound, the fact that the relative dualizing sheaf for C has degree $2g - 2$ gives the Lorenzini–Tucker bound. A similar argument was found independently by McCallum and Poonen [99].

We now explain where the subtlety occurs when one tries to combine the bounds of Stoll and Lorenzini–Tucker. As above, we define the Chabauty divisor

$$D_{\text{chab}} = \sum_{\bar{Q} \in \bar{C}^{\text{sm}}(\mathbb{F}_p)} n_{\bar{Q}}(\bar{Q})$$

and we let d be its degree. As in the case where C has good reduction, the goal is to show that $d \leq 2r$. When C has good reduction, Stoll proves this by combining the semicontinuity of h^0 and Clifford’s inequality. For singular curves, one can still define h^0 of a line bundle and it satisfies the desired semicontinuity theorem. However, even when C has semistable reduction, it is well-known that Clifford’s inequality does not hold in the form needed here. Katz and Zureick-Brown replace the use of Clifford’s inequality in Stoll’s argument by a hybrid between the classical Clifford inequality and Clifford’s inequality for linear series on tropical curves. In this way, they are able to obtain the desired bound $d \leq 2r$.

We briefly highlight the main steps in the argument, following the reformulation in terms of metrized complexes given in [5].

1. As noted by Katz and Zureick-Brown, if one makes a base change from \mathbb{Q}_p to an extension field over which there is a regular semistable model C' for C dominating the base change of C , the corresponding Chabauty divisors satisfy $D'_{\text{chab}} \geq D_{\text{chab}}$. We may therefore assume that C is a regular semistable model for C .
2. Let $s = \dim V_{\text{chab}} - 1$. We can identify V_{chab} with an $(s + 1)$ -dimensional space W of rational functions on C in the usual way by identifying $H^0(C, \Omega_C^1)$ with

$\mathcal{L}(K_C)$. The divisor D_{chab} on $\bar{\mathcal{C}}^{\text{sm}}$ defines in a natural way a divisor \mathcal{D} of degree d on the metrized complex \mathcal{C} associated with \mathcal{C} . We can promote the divisor $K_{\mathcal{C}} - \mathcal{D}$ to a limit linear series $(K_{\mathcal{C}} - \mathcal{D}, \{H_v\})$ by defining H_v to be the reduction of W to C_v for each $v \in V(G)$. By the definition of D_{chab} , each element of H_v vanishes to order at least $n_{\bar{Q}}$ at each point \bar{Q} in $\text{supp}(D_{\text{chab}}) \cap C_v$. The Specialization Theorem for limit linear series on metrized complexes then shows that

$$r_{\mathcal{C}}(K_{\mathcal{C}} - \mathcal{D}) \geq s \geq g - r - 1.$$

3. On the other hand, Clifford’s inequality for metrized complexes implies that

$$r_{\mathcal{C}}(K_{\mathcal{C}} - \mathcal{D}) \leq \frac{1}{2}(2g - 2 - d).$$

Combining these inequalities gives $d \leq 2r$ as desired.

11.2 The Uniformity Theorems of Katz–Rabinoff–Zureick-Brown

Together with Rabinoff, Katz, and Zureick-Brown have recently used linear series on tropical curves to refine another result due to Stoll. In [37], Caporaso, Harris, and Mazur proved that if one assumes the Bombieri–Lang conjecture then there is a uniform bound $M(g, K)$ depending only on g and the number field K such that $|C(K)| \leq M(g, K)$ for every curve C of genus $g \geq 2$ over K . The Bombieri–Lang conjecture, which asserts that the set of rational points on a variety of general type over a number field is not Zariski dense, remains wide open, and until recently little progress had been made in the direction of unconditional proofs of the Caporaso–Harris–Mazur result. In [114], Stoll proved that a uniform bound $M(g, K)$ exists for hyperelliptic curves provided that one assumes in addition that the Mordell–Weil rank of $\text{Jac}(C)(K)$ is at most $g - 3$. Katz, Rabinoff, and Zureick-Brown succeeded in removing the hypothesis in Stoll’s theorem that C is hyperelliptic, obtaining the following result.

Theorem 11.3 ([81]). *There is an explicit bound $N(g, d)$ such that if C is a curve of genus $g \geq 3$ defined over a number field K of degree d over \mathbb{Q} and having Mordell–Weil rank $r \leq g - 3$, then*

$$\#C(K) \leq N(g, d).$$

When $K = \mathbb{Q}$, one can take $N(g, 1) = 76g^2 - 82g + 22$.

Note that the bound in Theorem 11.2 is not uniform, because the quantity $|C^{\text{sm}}(\mathbb{F}_p)|$ can be arbitrarily large for a given prime p of bad reduction, and the smallest prime p of good reduction can be arbitrarily large as a function of g . Stoll’s main new idea was to apply the Chabauty–Coleman method on residue

annuli instead of just on discs. Stoll’s proof exploits the concrete description of differentials on a hyperelliptic curve as $f(x)dx/y$; the restriction of such a differential to an annulus has a bounded numerator, and Stoll is able to analyze the zeroes of the resulting p -adic integral via explicit computations with Newton polygons.

For general curves, such an explicit description of differentials and the Newton polygons of their p -adic integrals is not possible. This is where the theory of linear systems on metric graphs becomes useful. To circumvent the difficulty posed by not having an explicit description of differentials on C , Katz, Rabinoff, and Zureick-Brown generalize the Slope Formula (Theorem 6.4) to sections of a metrized line bundle. For a differential ω , the associated tropical function $F = \log |\omega|$ on the skeleton Γ of C belongs to the space $R(K_\Gamma^\#)$ of tropical rational functions G with $K_\Gamma^\# + \text{div}(G) \geq 0$. (The absolute value here comes from a natural formal metric on the canonical bundle.) Belonging to $R(K_\Gamma^\#)$ gives strong constraints on the slopes of F , and hence on the number of zeroes of the p -adic integral of ω . The Slope Formula thus replaces the Newton polygons in Stoll’s arguments, and estimates on the slopes of the Newton polygon are replaced by properties of the tropical linear series $|K_\Gamma^\#|$.

A major issue one faces in trying to establish Theorem 11.3 (which also shows up in the earlier work of Stoll) is that when C has bad reduction at p , there are two different kinds of p -adic integrals which need to be considered. On the one hand, there are the p -adic abelian integrals studied by Colmez, Zarhin, and Vologodsky, which have no periods and are obtained by pulling back the logarithm map on the p -adic Lie group $\text{Jac}(C)(\mathbb{Q}_p)$ to C . These are the integrals for which one knows that $\dim(V_{\text{chab}}) \geq g - r$. On the other hand, there are the p -adic integrals of Berkovich and Coleman–de Shalit which do have periods but also have better functoriality properties. These are the integrals which are given locally on residue annuli of a semistable model \mathcal{C} by formally integrating a local Laurent series expansion of $\omega \in H^0(C, \Omega^1)$. In order to prove Theorem 11.3, one needs to study the difference between the two kinds of p -adic integrals. One of the new discoveries of Katz, Rabinoff, and Zureick-Brown is that the difference can be understood quite concretely using tropical geometry by combining Theorem 6.5 with Raynaud’s uniformization theory.

The methods used by Katz–Rabinoff–Zureick-Brown in [81] also provide new results in the direction of a “uniform Manin-Mumford conjecture.” The Manin–Mumford conjecture, proved by Raynaud, asserts that if C is a curve of genus at least 2 embedded in its Jacobian via an Abel–Jacobi map $\iota : C \rightarrow \text{Jac}(C)$, then $\iota(C) \cap \text{Jac}(C)(\bar{K})_{\text{tors}}$ is finite. One can ask whether there is a uniform bound on the size of this intersection as one varies over all curves of a fixed genus g . The following uniform result for the number of K -rational points on C which are torsion on J is proved in [81]:

Theorem 11.4. *There is an explicit bound $N(g, d)_{\text{tors}}$ (which one can equal to the bound $N(g, d)$ above) such that if C is a curve of genus $g \geq 3$ defined over a number field K of degree d over \mathbb{Q} and $\iota : C \rightarrow \text{Jac}(C)$ is an Abel–Jacobi embedding defined over K , then*

$$\#\iota(C) \cap \text{Jac}(C)(K)_{\text{tors}} \leq N(g, d)_{\text{tors}}.$$

Note that in Theorem 11.4 there is no restriction on the Mordell–Weil rank of $\text{Jac}(C)$. It has been conjectured that $\#A(K)_{\text{tors}}$ is bounded uniformly in terms of $[K : \mathbb{Q}]$ and g for all abelian varieties of dimension g over K , which would of course imply Theorem 11.4 as a special case, but this is known only for $g = 1$ [100] and the general case seems far out of reach at present.

Katz, Rabinoff, and Zureick-Brown also prove a uniformity result concerning the number of *geometric* (\bar{K} -rational) torsion points lying on C , under a technical assumption about the structure of the stable model at some prime p . We refer to [81] for the precise statement.

12 Limiting Behavior of Weierstrass Points in Degenerating Families

The theory discussed in this paper has interesting applications to the behavior of Weierstrass points under specialization. To motivate this kind of question, we begin with a seemingly unrelated classical result due to Andrew Ogg [106].

12.1 Weierstrass Points on Modular Curves

Let N be a positive integer. The finite-dimensional space $S = S_2(\Gamma_0(N))$ of weight 2 cusp forms for the congruence subgroup $\Gamma_0(N)$ of $\text{SL}_2(\mathbb{Z})$ is an important object in number theory. An element $f \in S$ has a *q-expansion* of the form $f = \sum_{n=1}^{\infty} a_n q^n$ with $a_n \in \mathbb{C}$, which uniquely determines f . For $f \neq 0$ in S , define

$$\text{ord}(f) = \inf\{n \mid a_n \neq 0\} - 1. \tag{4}$$

If $g = g_0(N) = \dim(S)$, then by Gaussian elimination there exists an element $f \in S$ with $\text{ord}(f) \geq g - 1$. Is there any unexpected cancellation? Under certain restrictions on the level N , the answer is no:

Theorem 12.1 ([106]). *If $N = pM$ with p prime, $p \nmid M$, and $g_0(M) = 0$, then there is no nonzero element f of $S_2(\Gamma_0(N))$ with $\text{ord}(f) \geq g$. (In particular, this holds if $N = p$ is prime.)*

One can give an enlightening proof of Ogg’s theorem using specialization of divisors from curves to metric graphs; the following argument is taken from [17].

First of all, Ogg’s theorem can be recast in the following purely geometric way, which is in fact how Ogg formulated and proved the result in [106]:

Theorem 12.2. *If $N = pM$ with p prime, $p \nmid M$, and $g_0(M) = 0$, then the cusp ∞ is not a Weierstrass point on the modular curve $X_0(N)$.*

Recall that a point P on a genus g curve X is called a *Weierstrass point* if there exists a holomorphic differential $\omega \in H^0(X, \Omega_X^1)$ vanishing to order at least g at P . To see the equivalence between Theorems 12.1 and 12.2, recall that q is an analytic local parameter on $X_0(N)$ at the cusp ∞ and the map $f \mapsto f(q)\frac{dq}{q}$ gives an isomorphism between $S_2(\Gamma_0(N))$ and the space of holomorphic differentials on $X_0(N)$. Under this isomorphism, the function ord defined in (4) becomes the order of vanishing of the corresponding differential at ∞ . So there is a nonzero element f of $S_2(\Gamma_0(N))$ with $\text{ord}(f) \geq g$ if and only if there is a nonzero holomorphic differential ω vanishing to order at least g at ∞ .


The reduction of $X = X_0(N)$ modulo p when p exactly divides $N = pM$ is well-understood; the special fiber of the so-called *Deligne–Rapoport model* for $X_0(N)$ over \mathbb{Z}_p consists of two copies of $X_0(M)$ intersecting transversely at the supersingular points in characteristic p . This model is always semistable but is not in general regular. (It is very easy to describe the minimal regular model, but we will not need this here.) In any case, the skeleton Γ of $X_0(N)$ over \mathbb{Q}_p is a “metric banana graph” consisting of two vertices connected by a number of edges, as pictured in Fig. 16, and the cusp ∞ specializes to one of the two vertices, call it P . Under the hypotheses of Theorem 12.1, each $X_0(M)$ is a rational curve and so the genus of Γ is equal to the genus of $X_0(N)$. That is, there are $g + 1$ edges. By the Specialization Theorem, if there is a nonzero global section of K_X vanishing to order at least g at ∞ , then $r(K_\Gamma - gP) \geq 0$. However, since $K_\Gamma = (g - 1)P + (g - 1)Q$, where Q is the other vertex, we have $K_\Gamma - gP = (g - 1)Q - P$, which is P -reduced by Dhar’s algorithm and non-effective. Therefore $r(K_\Gamma - gP) = -1$, and Ogg’s theorem is proved.

12.2 Specialization of Weierstrass Points

The essence of the above argument is that if C is a *totally degenerate* curve, meaning that the genus of its minimal skeleton Γ equals the genus of C , then the Weierstrass points on C must specialize to Weierstrass points on Γ , where a Weierstrass point on Γ is a point P such that $r(K_\Gamma - gP) \geq 0$. It follows from the Specialization Theorem and the corresponding fact from algebraic geometry that if Γ is a metric graph of genus $g \geq 2$ then the set of Weierstrass points on Γ is nonempty. A purely combinatorial proof of this fact was given by Amini [3].

The specialization of Weierstrass points is also a natural thing to study from the purely algebro-geometric point of view, where one is asking about the limiting

Fig. 16 The “Banana” graph of genus 2



behavior of the Weierstrass points in a semistable one-parameter family of curves. This subject, which was previously studied by Eisenbud–Harris [59], Esteves–Medeiros [60], and several other authors, has seen important recent advances by Amini [4]. We now summarize the main results proved in Amini’s paper.

Let L be a line bundle of degree d and rank $r \geq 0$ on a curve C of genus g over an algebraically closed field k of characteristic zero. Given a point $P \in C(k)$, we define the *vanishing set* $S_P(L)$ of L at P to be the set of orders of vanishing of global sections of L at P . We have $|S_P(L)| = r + 1$ for all $P \in C(k)$, and for all but finitely many $P \in C(k)$ the vanishing set is $[r] := \{0, 1, \dots, r\}$. A point $P \in C(k)$ whose vanishing set is not $[r]$ is called a *Weierstrass point* for L . Equivalently, P is a Weierstrass point for L if there exists a global section of L vanishing to order at least $r + 1$ at P . A Weierstrass point of C is by definition a Weierstrass point for the canonical bundle K_C .

The L -weight of a point $P \in C(k)$ is

$$\text{wt}_P(L) = \left(\sum_{m \in S_P(L)} m \right) - \binom{r+1}{2} = \sum_{m \in S_P(L)} m - \sum_{i \in [r]} i.$$

Thus $\text{wt}_P(L) \geq 0$ for all $P \in C(k)$ and $\text{wt}_P(L) > 0$ if and only if P is a Weierstrass point for L . The *Weierstrass divisor* for L is $\mathcal{W} = \mathcal{W}(L) = \sum_{P \in C(k)} \text{wt}_P(L)(P)$. If we fix a basis \mathcal{F} for $H^0(C, L)$, the corresponding *Wronskian* $\text{Wr}_{\mathcal{F}}$ is a nonzero global section of $L^{\otimes(r+1)} \otimes K_C^{\otimes \frac{r(r+1)}{2}}$ whose divisor is precisely $\mathcal{W}(L)$. In particular, the degree of $\mathcal{W}(L)$ (i.e., the total number of Weierstrass points counted according to their weights) is $W(L) := d(r + 1) + (g - 1)r(r + 1)$.

We seek an explicit formula for $\text{Trop}(\mathcal{W})$. For this, it is convenient to fix a divisor with $L = L(D)$, and to define as usual $\mathcal{L}(D) = \{f \in k(C)^* \mid \text{div}(f) + D \geq 0\}$. Let $D_\Gamma = \sum_{x \in \Gamma} d_x(x)$ be the specialization $\text{Trop}(D)$ of D to Γ . Let $K_\Gamma^\#$ be the canonical divisor of Γ considered as a vertex-weighted metric graph, as in Sect. 4.4. Concretely, we have $K_\Gamma^\# = \sum_{x \in \Gamma} (2g_x - 2 + \text{val}(x))x$.

For a tangent direction ν at x , define $S^\nu(D)$ to be the set of integers occurring as $s^\nu(f)$ for some $f \in \mathcal{L}(D)$, where $s^\nu(f)$ is defined as in Sect. 6 to be the slope of $\text{trop}(f)$ in the tangent direction ν . Since $s^\nu(f)$ coincides with the order of vanishing of the normalized reduction \tilde{f}_x at the point of C_x corresponding to ν , one sees easily that $|S^\nu(D)| = r + 1$.

For $x \in \Gamma$, let

$$S_x(D) = \begin{cases} \sum_{\nu \in T_x(\Gamma)} \left(\sum_{s \in S^\nu(D)} s \right) & \text{if } x \text{ is of type-2} \\ 0 & \text{otherwise,} \end{cases}$$

where $T_x(\Gamma)$ denotes the set of tangent directions at x in Γ , and let

$$S(D) = \sum_{x \in \Gamma} S_x(D) x.$$

Note that $\deg(S(D)) = 0$, since if $f \in k(C)^*$ then the slope of $F = -\log |f|$ along an oriented edge \vec{e} of Γ is the negative of the slope of F along the same edge with the orientation reversed.

The following formula is due to Amini. When Γ is the skeleton of a semistable R -model \mathcal{C} for C , the formula shows how the Weierstrass points of the generic fiber C specialize to the various components of the special fiber of \mathcal{C} , providing a simple and satisfying answer to a question of Eisenbud and Harris.

Theorem 12.3 ([3]). *Let $\text{Trop} : \text{Div}(C_{\bar{k}}) \rightarrow \text{Div}(\Gamma)$ be the natural map. Then*

$$\text{Trop}(\mathcal{W}(L)) = (r + 1) \text{Trop}(D) + \binom{r + 1}{2} K_{\Gamma}^{\#} - S(D). \tag{5}$$

Note that since $\deg(S(D)) = 0$, the degree of the right-hand side of (5) is $W(L) = \deg(\mathcal{W}(L))$ as expected. Amini also proves an analogue of (5) when the residue field of k has positive characteristic. As this is more technical to state, we will not discuss this here.

Remark 12.4. A metric graph can have infinitely many Weierstrass points; this happens, for example, with the banana graphs of genus $g \geq 3$ discussed above (see [17]). In general, the set of Weierstrass points on a metric graph Γ is a finite disjoint union of closed connected sets. It is an open problem to determine whether there are intrinsic multiplicities $m(A)$ attached to each connected component A of the Weierstrass locus on a metric graph Γ such that for any curve C having Γ as a skeleton, exactly $m(A)$ Weierstrass points of C tropicalize to A .

12.3 Distribution of Weierstrass Points

Amini uses formula (5) to prove a non-Archimedean analogue of the Mumford–Neeman equidistribution theorem, previously conjectured by Baker. We first recall the statement of the latter result, and then present Amini’s analogous theorem.

Let C be a compact Riemann surface of genus at least 1. There is a natural volume form ω_{Ar} on C , called the *Arakelov form*, which can be defined as follows. Let $\omega_1, \dots, \omega_g$ be an orthonormal basis of $\mathcal{L}(K_C)$ with respect to the Hermitian inner product

$$\langle \omega, \nu \rangle = \frac{i}{2} \int_C \omega \wedge \bar{\nu}.$$

Then the $(1, 1)$ -form $\omega_C = \frac{i}{2} \sum_{j=1}^g \omega_j \wedge \bar{\omega}_j$ does not depend on the choice of $\omega_1, \dots, \omega_g$ and has total mass g . We define

$$\omega_{Ar} := \frac{1}{g} \omega_C.$$

Geometrically, the curvature form of ω_C is the pullback of the curvature form of the flat metric on the Jacobian J of C with respect to any Abel–Jacobi map $C \rightarrow J$. Since the flat metric on J is translation-invariant, the pullback in question is independent of the choice of base point in the definition of the Abel–Jacobi map.

The Mumford–Neeman theorem [103] asserts that for any ample line bundle L on C , the Weierstrass points of $L^{\otimes n}$ become equidistributed with respect to ω_{Ar} as n tends to infinity:

Theorem 12.5. *Let C be a compact Riemann surface of genus at least 1 and let L be an ample line bundle on C . Let*

$$\delta_n = \frac{1}{W(L^{\otimes n})} \sum_{P \in C} \text{wt}_P(L^{\otimes n}) \delta_P$$

be the probability measure supported equally on the Weierstrass points of $L^{\otimes n}$. Then as n tends to infinity, the measures δ_n converge weakly¹⁵ to the Arakelov metric ω_{Ar} .

In order to state Amini’s non-Archimedean analogue of Theorem 12.5, we will first define the analogue of the Zhang measure on vertex-weighted metric graphs/Berkovich curves following [18].

Let Γ be a metric graph of genus g . We fix a weighted graph model G of Γ and for each edge e of G let $\ell(e)$ denote the length of e . For each spanning tree T of G , let e_1, \dots, e_g denote the edges of G not belonging to T , and let

$$\mu_T = \sum_{j=1}^g \lambda(e_j)$$

where $\lambda(e)$ is Lebesgue measure along e , normalized to have total mass 1 (so that μ_T has total mass g). We also let $w(T) = \prod_{j=1}^g \ell(e_j)$, and let

$$w(G) = \sum_T w(T)$$

be the sum of $w(T)$ over all spanning trees T of G . Then the measure

$$\mu_\Gamma = \sum_T \frac{w(T)}{w(G)} \mu_T$$

is a weighted average of the measures μ_T over all spanning trees T , and in particular has total mass g .

¹⁵This means that for every continuous function $f : C \rightarrow \mathbb{R}$, we have $\int_C f \delta_n = \int_C f \omega_{Ar}$.

In other words, a random point in the complement of a random spanning tree of G is distributed according to the probability measure $\frac{1}{g}\mu_\Gamma$.

Now let (Γ, ω) be a vertex-weighted metric graph, in the sense of Sect. 4.4, of genus $g = g(\Gamma) + \sum_{x \in \Gamma} \omega(x)$. Then the measure

$$\mu_{(\Gamma, \omega)} := \mu_\Gamma + \sum_{x \in \Gamma} \omega_x \delta_x$$

has total mass g . If $g \geq 1$, we define the *Zhang measure* on (Γ, ω) to be the probability measure

$$\mu_{\text{Zh}} := \frac{1}{g} \mu_{(\Gamma, \omega)}.$$

Theorem 12.6 ([4]). *Let C be an algebraic curve of genus at least 1 over the non-Archimedean field k of equal characteristic 0, and let L be an ample line bundle on C . Let \mathcal{C} be a strongly semistable model of C over the valuation ring of k , let (Γ, ω) be the weighted graph associated with \mathcal{C} in the sense of Sect. 4.4, and let μ_{Zh} be the Zhang measure associated with (Γ, ω) . Finally, let*

$$\delta_n = \frac{1}{W(L^{\otimes n})} \sum_{P \in \mathcal{C}} \text{wt}_P(L^{\otimes n}) \delta_{\text{Trop}(P)}$$

be the probability measure on Γ supported equally on the tropicalizations of the Weierstrass points of $L^{\otimes n}$ (taken with multiplicities). Then as n tends to infinity, the measures δ_n converge weakly on Γ to μ_{Zh} .

A new and concrete consequence of Theorem 12.6 is the following:

Corollary 12.7. *Let C be an algebraic curve of genus at least 1 over a non-Archimedean field k of equal characteristic 0, and let L be an ample line bundle on C . Fix a strongly semistable model \mathcal{C} for C over the valuation ring of k , let Z be an irreducible component of the special fiber of \mathcal{C} , and let g_Z be the genus of Z . Let $\mathcal{W}_Z(L^{\otimes n})$ be the set of Weierstrass points of $L^{\otimes n}$ specializing to a nonsingular point of Z . Then*

$$\lim_{n \rightarrow \infty} \frac{|\mathcal{W}_Z(L^{\otimes n})|}{|\mathcal{W}(L^{\otimes n})|} = \frac{g_Z}{g}.$$

It is convenient and enlightening to rephrase Theorem 12.6 in terms of the Berkovich analytic space C^{an} . If Γ is any skeleton of C^{an} and ω is the corresponding weight function defined by $\omega(x) = g_x$, we define the *Zhang measure* on C^{an} to be the probability measure

$$\mu_{\text{Zh}} := \frac{1}{g} \iota_* \mu_{(\Gamma, \omega)}$$

with respect to the natural inclusion $\iota : \Gamma \rightarrow C^{\text{an}}$. This measure on C^{an} is easily seen to be independent of the choice of Γ , and has total mass equal to the genus g of C . Using the fact that $C^{\text{an}} \cong \varprojlim \Gamma_C$ (cf. Sect. 6.2), one deduces using standard results from real analysis that Theorem 12.6 is equivalent to the following reformulation, which more closely resembles Theorem 12.5:

Theorem 12.8 ([4]). *Let C be an algebraic curve of genus at least 1 over the non-Archimedean field k of equal characteristic 0, and let L be an ample line bundle on C . Let*

$$\delta_n = \frac{1}{W(L^{\otimes n})} \sum_{P \in C} \text{wt}_P(L^{\otimes n}) \delta_P$$

be the probability measure on C^{an} supported equally on the Weierstrass points of $L^{\otimes n}$. Then as n tends to infinity, the measures δ_n converge weakly on C^{an} to the Zhang measure μ_{Zh} .

Remark 12.9. The measure μ_{Zh} , which was first introduced in the context of vertex-weighted metric graphs by Shouwu Zhang in [120], plays the role in the non-Archimedean setting of the Arakelov volume form. Using a result of Heinz [74] and the recent work of Chambert-Loir–Ducros [44] and Gubler–Kunemann [69] on non-Archimedean Arakelov theory, one can show that, as in the Archimedean case, μ_C is obtained by pulling back the curvature form of a canonical translation-invariant metric on J via an Abel–Jacobi map. There is also evidently a close connection between the measure μ_Γ and the polyhedral decomposition $\{C_T\}$ of $\text{Pic}^g(\Gamma)$ associated with G (cf. Sect. 5.4) which is deserving of further study.

The proof of Theorem 12.6 (and its equivalent formulation Theorem 12.8) is based on formula (5) together with the theory of *Okounkov bodies*. The rough idea is that fixing a type-2 point x of C^{an} and a tangent direction v at x , as well as a divisor D with $L = L(D)$, the rational numbers $\frac{1}{n}S^v(nD)$ defined above become equidistributed in a real interval of length $d = \text{deg}(L)$ as $n \rightarrow \infty$. Combining this “local” equidistribution result with (5) and a careful analysis of the variation of the minimum slope along edges of Γ give the desired result.

13 Further Reading

There are many topics closely related to the contents of this paper which we have not had space to discuss. Here is a brief and non-exhaustive list of some related topics and papers which we recommend to the interested reader:

1. *Harmonic morphisms.* In algebraic geometry, a base-point free linear series of rank r on a curve C is more or less the same thing as a morphism $C \rightarrow \mathbb{P}^r$. In tropical geometry, the situation is much more subtle, and no satisfactory analogue of this correspondence is known. For $r = 1$, there is a close relationship

(although not a precise correspondence) between tropical g_d^1 's on a metric graph Γ and degree d harmonic morphisms from Γ to a metric tree. The theory of harmonic morphisms of metric graphs and metrized complexes of curves is explored in detail in the papers [8, 9, 45, 97], among others.

2. *Spectral bounds for gonality.* In [49], Cornelissen et. al. establish a spectral lower bound for the *stable gonality* (in the sense of harmonic morphisms) of a graph G in terms of the smallest nonzero eigenvalue of the Laplacian of G . This is a tropical analogue of the Li–Yau inequality for Riemann surfaces. They give applications of their tropical Li–Yau inequality to uniform boundedness of torsion points on rank two Drinfeld modules, as well as to lower bounds on the modular degree of elliptic curves over function fields. The spectral bound from [49] was subsequently refined by Amini and Kool in [7] to a spectral lower bound for the *divisorial gonality* (i.e., the minimal degree of a rank 1 divisor) of a metric graph Γ . In [7], as well as in the related paper [54], this circle of ideas is applied to show that the expected gonality of a random graph is asymptotic to the number of vertices.
3. *Tropical complexes.* In [39], Cartwright formulates a higher-dimensional analogue of the basic theory of linear series on graphs, including a Specialization Theorem for the rank function. He calls the objects on which his higher-dimensional linear series live *tropical complexes*. A generalization of the Slope Formula to the context of non-Archimedean varieties and tropical complexes is proved in [70].
4. *Abstract versus embedded tropical curves.* In this paper we have dealt exclusively with linear series on abstract tropical curves (thought of as metric graphs) and have eschewed the more traditional perspective of tropical varieties as non-Archimedean amoebas associated with subvarieties of tori. The two approaches are closely related, however: see, for example, [24, 42, 70]. The theory of linear series on abstract tropical curves has concrete consequences for embedded tropical curves, e.g., with respect to the theory of bitangents and theta characteristics as in [26, 46].
5. *Algebraic rank.* In [35], Caporaso introduces a notion of rank for divisors on graphs known as the *algebraic rank*, which is defined geometrically by varying over all curves with the given dual graph and all line bundles with the given specialization. The algebraic rank differs in general from the combinatorial rank [38], but the two invariants agree for hyperelliptic graphs and graphs of genus 3 [82]. Many of the results we have discussed also hold for the algebraic rank. For example, there are specialization, Riemann–Roch, and Clifford's theorems for algebraic rank [35], and Mnëv universality holds for obstructions to the lifting problem for algebraic rank [91].

Acknowledgements The authors would like to thank Omid Amini and Sam Payne for enlightening discussions, and Eric Katz and Joe Rabinoff for helpful feedback on our summary of Katz et al. [81]. They also thank Spencer Backman, Dustin Cartwright, Melody Chan, Yoav Len, and Sam Payne for their comments on an early draft of this paper.

References

1. Abramovich, D., Caporaso, L., Payne, S.: The tropicalization of the moduli space of curves. *Ann. Sci. Éc. Norm. Supér.* **48**(4), 765–809 (2015)
2. Agrawal, R., Musiker, G., Sotirov, V., Wei, F.: Involutions on standard Young tableaux and divisors on metric graphs. *Electron. J. Comb.* **20**(3) (2013)
3. Amini, O.: Reduced divisors and embeddings of tropical curves. *Trans. Am. Math. Soc.* **365**(9), 4851–4880 (2013)
4. Amini, O.: Equidistribution of Weierstrass points on curves over non-Archimedean fields. Preprint, arXiv:1412.0926v1 (2014)
5. Amini, O., Baker, M.: Linear series on metrized complexes of algebraic curves. *Math. Ann.* **362**(1–2), 55–106 (2015)
6. Amini, O., Caporaso, L.: Riemann–Roch theory for weighted graphs and tropical curves. *Adv. Math.* **240**, 1–23 (2013)
7. Amini, O., Kool, F.: A spectral lower bound for the divisorial gonality of metric graphs. Preprint, arXiv:1407.5614 (2014)
8. Amini, O., Baker, M., Brugallé, E., Rabinoff, J.: Lifting harmonic morphisms I: metrized complexes and Berkovich skeleta. *Res. Math. Sci.* **2**, Art. 7, 67 (2015)
9. Amini, O., Baker, M., Brugallé, E., Rabinoff, J.: Lifting harmonic morphisms II: tropical curves and metrized complexes. *Algebra Number Theory.* **9**(2), 267–315 (2015)
10. An, Y., Baker, M., Kuperberg, G., Shokrieh, F.: Canonical representatives for divisor classes on tropical curves and the matrix–tree theorem. *Forum Math. Sigma* **2**, e24, 25 (2014)
11. Arbarello, E., Ciliberto, C.: Adjoint hypersurfaces to curves in \mathbf{P}^r following Petri. In: *Commutative Algebra (Trento, 1981)*. Lecture Notes in Pure and Applied Mathematics, vol. 84, pp. 1–21. Dekker, New York (1983)
12. Arbarello, E., Cornalba, M., Griffiths, P.A., Harris, J.: *Geometry of Algebraic Curves. Volume I*. Grundlehren der Mathematischen Wissenschaften, vol. 267. Springer, New York (1985)
13. Babbage, D.W.: A note on the quadrics through a canonical curve. *J. Lond. Math. Soc.* **14**, 310–315 (1939)
14. Backman, S.: Infinite reduction of divisors on metric graphs. *European J. Comb.* **35**, 67–74 (2014)
15. Backman, S.: Riemann–Roch theory for graph orientations. Preprint, arXiv:1401.3309 (2014)
16. Baker, M.: An introduction to Berkovich analytic spaces and non-Archimedean potential theory on curves. In: *p -Adic Geometry*. University Lecture Series, vol. 45, pp. 123–174. American Mathematical Society, Providence, RI (2008)
17. Baker, M.: Specialization of linear systems from curves to graphs. *Algebra Number Theory* **2**(6), 613–653 (2008)
18. Baker, M., Faber, X.: Metric properties of the tropical Abel–Jacobi map. *J. Algebraic Combin.* **33**(3), 349–381 (2011)
19. Baker, M., Norine, S.: Riemann–Roch and Abel–Jacobi theory on a finite graph. *Adv. Math.* **215**(2), 766–788 (2007)
20. Baker, M., Rabinoff, J.: The skeleton of the Jacobian, the Jacobian of the skeleton, and lifting meromorphic functions from tropical to algebraic curves. *Int. Math. Res. Not.* **2015**(16), 7436–7472 (2015)
21. Baker, M., Rumely, R.: Potential theory and dynamics on the Berkovich projective line. *Mathematical Surveys and Monographs*, vol. 159. American Mathematical Society, Providence, RI (2010)
22. Baker, M., Shokrieh, F.: Chip-firing games, potential theory on graphs, and spanning trees. *J. Comb. Theory Ser. A* **120**(1), 164–182 (2013)
23. Baker, M., Wang, Y.: The Bernardi process and torsor structures on spanning trees. Preprint, arXiv:1406.1584 (2014)
24. Baker, M., Payne, S., Rabinoff, J.: Nonarchimedean geometry, tropicalization, and metrics on curves. *Algebraic Geom.* Preprint, arXiv:1104.0320v3 (2011, to appear)

25. Baker, M., Payne, S., Rabinoff, J.: On the structure of non-Archimedean analytic curves. In: Tropical and Non-Archimedean Geometry. Contemporary Mathematics, vol. 605, pp. 93–121. American Mathematical Society, Providence, RI (2013)
26. Baker, M., Len, Y., Morrison, R., Pflueger, N., Ren, Q.: Bitangents of tropical plane quartic curves. Preprint, arXiv:1404.7568 (2014)
27. Ballico, E., Ellia, P.: The maximal rank conjecture for nonspecial curves in \mathbf{P}^3 . *Invent. Math.* **79**(3), 541–555 (1985)
28. Bayer, D., Eisenbud, D.: Graph curves. *Adv. Math.* **86**(1), 1–40 (1991) With an appendix by Sung Won Park.
29. Berkovich, V.: Spectral Theory and Analytic Geometry over Non-Archimedean Fields. Mathematical Surveys and Monographs, vol. 33. American Mathematical Society, Providence, RI (1990)
30. Biggs, N.L.: Chip-firing and the critical group of a graph. *J. Algebraic Combin.* **9**(1), 25–45 (1999)
31. Brannetti, S., Melo, M., Viviani, F.: On the tropical Torelli map. *Adv. Math.* **226**(3), 2546–2586 (2011)
32. Caporaso, L.: Néron models and compactified Picard schemes over the moduli stack of stable curves. *Am. J. Math.* **130**(1), 1–47 (2008)
33. Caporaso, L.: Algebraic and tropical curves: comparing their moduli spaces. In: Farkas and Morrison (eds.) Handbook of Moduli. Adv. Lect. Math. (ALM), vol. 24, pp. 119–160. Int. Press, Somerville (2013)
34. Caporaso, L.: Algebraic and combinatorial Brill-Noether theory. In: Compact Moduli Spaces and Vector Bundles. Contemporary Mathematics, vol. 564, pp. 69–85. American Mathematical Society, Providence, RI (2012)
35. Caporaso, L.: Rank of divisors on graphs: an algebro-geometric analysis. In: A Celebration of Algebraic Geometry. Clay Mathematics Proceedings, vol. 18, pp. 45–64. American Mathematical Society, Providence, RI (2013)
36. Caporaso, L., Viviani, F.: Torelli theorem for graphs and tropical curves. *Duke Math. J.* **153**(1), 129–171 (2010)
37. Caporaso, L., Harris, J., Mazur, B.: Uniformity of rational points. *J. Am. Math. Soc.* **10**(1), 1–35 (1997)
38. Caporaso, L., Len, Y., Melo, M.: Algebraic and combinatorial rank of divisors on finite graphs. Preprint, arXiv:1401:5730v2 (2014)
39. Cartwright, D.: Tropical complexes. Preprint, arXiv:1308.3813 (2013)
40. Cartwright, D.: Lifting matroid divisors on tropical curves. Preprint, arXiv:1502.03759 (2015)
41. Cartwright, D., Jensen, D., Payne, S.: Lifting divisors on a generic chain of loops. *Canad. Math. Bull.* **58**(2), 250–262 (2015)
42. Cartwright, D., Dudzik, A., Manjunath, M., Yao, Y.: Embeddings and immersions of tropical curves. Preprint, arXiv:1409.7372v1 (2014)
43. Castelnuovo, G., Enriques, F., Severi, F.: Max Noether. *Math. Ann.* **93**(1), 161–181 (1925)
44. Chambert-Loir, A., Ducros, A.: Formes différentielles réelles et courants sur les espaces de Berkovich. Preprint, arXiv:1204.6277 (2012)
45. Chan, M.: Tropical hyperelliptic curves. *J. Algebraic Combin.* **37**(2), 331–359 (2013)
46. Chan, M., Jiradilok, P.: Theta characteristics of tropical K_4 curves. Preprint, arXiv:1503:05776v1 (2015)
47. Ciliberto, C., Harris, J., Miranda, R.: On the surjectivity of the Wahl map. *Duke Math. J.* **57**(3), 829–858 (1988)
48. Cools, F., Draisma, J., Payne, S., Robeva, E.: A tropical proof of the Brill-Noether theorem. *Adv. Math.* **230**(2), 759–776 (2012)
49. Cornelissen, G., Fumiharu, K., Kool, J.: A combinatorial Li–Yau inequality and rational points on curves. *Math. Ann.* **361**(1–2), 211–258 (2015)
50. Coleman, R.F.: Effective Chabauty. *Duke Math. J.* **52**(3), 765–770 (1985)
51. Conrad, B.: Several approaches to non-Archimedean geometry. In: p -Adic Geometry. University Lecture Series, vol. 45, pp. 9–63. American Mathematical Society, Providence, RI (2008)

52. Coppens, M.: Clifford's theorem for graphs. Preprint, arXiv:1304.6101 (2014)
53. Coppens, M.: Free divisors on metric graphs. Preprint, arXiv:1410.3114 (2014)
54. Deveau, A., Jensen, D., Kainic, J., Mitropolsky, D.: Gonality of random graphs. Preprint, arXiv:1409.5688 (2014)
55. Dhar, D.: Self-organized critical state of sandpile automaton models. *Phys. Rev. Lett.* **64**(14), 1613–1616 (1990)
56. Dhar, D., Ruelle, P., Sen, S., Verma, D.-N.: Algebraic aspects of Abelian sandpile models. *J. Phys. A* **28**(4), 805–831 (1995)
57. Eisenbud, D., Harris, J.: A simpler proof of the Gieseker-Petri theorem on special divisors. *Invent. Math.* **74**(2), 269–280 (1983)
58. Eisenbud, D., Harris, J.: Limit linear series: basic theory. *Invent. Math.* **85**(2), 337–371 (1986)
59. Eisenbud, D., Harris, J.: Existence, decomposition, and limits of certain Weierstrass points. *Invent. Math.* **87**(3), 495–515 (1987)
60. Esteves, E., Medeiros, N.: Limits of Weierstrass points in regular smoothings of curves with two components. *C. R. Acad. Sci. Paris Sér. I Math.* **330**(10), 873–878 (2000)
61. Farkas, G.: Koszul divisors on moduli spaces of curves. *Am. J. Math.* **131**(3), 819–867 (2009)
62. Foster, T., Gross, P., Payne, S.: Limits of tropicalizations. *Israel J. Math.* **201**(2), 835–846 (2014)
63. Gathmann, A., Kerber, M.: A Riemann-Roch theorem in tropical geometry. *Math. Z.* **259**(1), 217–230 (2008)
64. Gathmann, A., Kerber, M., Markwig, H.: Tropical fans and the moduli spaces of tropical curves. *Compos. Math.* **145**(1), 173–195 (2009)
65. Gieseker, D.: Stable curves and special divisors: Petri's conjecture. *Invent. Math.* **66**(2), 251–275 (1982)
66. Griffiths, P., Harris, J.: On the variety of special linear systems on a general algebraic curve. *Duke Math. J.* **47**(1), 233–272 (1980)
67. Gubler, W.: Tropical varieties for non-Archimedean analytic spaces. *Invent. Math.* **169**(2), 321–376 (2007)
68. Gubler, W.: Non-Archimedean canonical measures on abelian varieties. *Compos. Math.* **146**(3), 683–730 (2010)
69. Gubler, W., Künnemann, K.: A tropical approach to non-Archimedean Arakelov theory. Preprint, arXiv:1406.7637 (2014)
70. Gubler, W., Rabinoff, J., Werner, A.: Skeletons and tropicalizations. Preprint, arXiv:1404.7044 (2014)
71. Harris, J.: *Curves in Projective Space*. Séminaire de Mathématiques Supérieures, vol. 85 Presses de l'Université de Montréal, Montreal, Que (1982) With the collaboration of David Eisenbud
72. Harris, J., Morrison, I.: *Moduli of Curves*. Graduate Texts in Mathematics, vol. 187. Springer, New York (1998)
73. Haase, C., Musiker, G., Yu, J.: Linear systems on tropical curves. *Math. Z.* **270**(3–4), 1111–1140 (2012)
74. Heinz, N.: Admissible metrics for line bundles on curves and abelian varieties over non-Archimedean local fields. *Arch. Math. (Basel)* **82**(2), 128–139 (2004)
75. Hladký, J., Král', D., Norine, S.: Rank of divisors on tropical curves. *J. Combin. Theory Ser. A* **120**(7), 1521–1538 (2013)
76. Holroyd, A., Levine, L., Mészáros, K., Peres, Y., Propp, J., Wilson, D.: Chip-firing and rotor-routing on directed graphs. In: *In and Out of Equilibrium 2*. Progress in Probability, vol. 60, pp. 331–364. Birkhäuser, Basel (2008)
77. Jensen, D.: The locus of Brill-Noether general graphs is not dense. Preprint, arXiv:1405.6338 (2014)
78. Jensen, D., Payne, S.: Tropical independence I: Shapes of divisors and a proof of the Gieseker-Petri theorem. *Algebra Number Theory* **8**(9), 2043–2066 (2014)
79. Jensen, D., Payne, S.: Tropical independence II: the maximal rank conjecture for quadrics. *Algebra Number Theory*. <http://arxiv.org/abs/1505.05460>

80. Katz, E., Zureick-Brown, D.: The Chabauty-Coleman bound at a prime of bad reduction and Clifford bounds for geometric rank functions. *Compos. Math.* **149**(11), 1818–1838 (2013)
81. Katz, E., Rabinoff, J., Zureick-Brown, D.: Uniform bounds for the number of rational points on curves of small Mordell-Weil rank. *Duke Math J.* arXiv:1504.00694 (2015)
82. Kawaguchi, S., Yamaki, K.: Algebraic rank of hyperelliptic graphs and graphs of genus 3. Preprint, arXiv:1401:3935 (2014)
83. Kawaguchi, S., Yamaki, K.: Rank of divisors on hyperelliptic curves and graphs under specialization. Preprint, arXiv:1304.6979v3 (2014)
84. Kempf, G.: *Schubert Methods with an Application to Algebraic Curves*. Publ. Math. Centrum (1971)
85. Kiss, V., Tóthmérész, L.: Chip-firing games on Eulerian digraphs and NP-hardness of computing the rank of a divisor on a graph. Preprint, arXiv:1407.6598 (2014)
86. Kleiman, S.L., Laksov, D.: Another proof of the existence of special divisors. *Acta Math.* **132**, 163–176 (1974)
87. Kozlov, D.: The topology of moduli spaces of tropical curves with marked points. *Asian J. Math.* **13**(3), 385–403 (2009)
88. Lafforgue, L.: *Chirurgie des grassmanniennes*. CRM Monograph Series. vol. 19. American Mathematical Society, Providence, RI (2003)
89. Lazarsfeld, R.: Brill-Noether-Petri without degenerations. *J. Differential Geom.* **23**(3), 299–307 (1986)
90. Len, Y.: The Brill-Noether rank of a tropical curve. Preprint, arXiv:1209.6309v1 (2012)
91. Len, Y.L. A note on algebraic rank, matroids, and metrized complexes. Preprint, arXiv:1401:8156 (2014)
92. Levine, L., Propp, J.: What is . . . a sandpile? *Not. Am. Math. Soc.* **57**(8), 976–979 (2010)
93. Lim, C.-M., Payne, S., Potashnik, N.: A note on Brill-Noether theory and rank-determining sets for metric graphs. *Int. Math. Res. Not. IMRN* **23**:5484–5504 (2012)
94. Liu, Q.: *Algebraic Geometry and Arithmetic Curves*. Oxford Graduate Texts in Mathematics, vol. 6. Oxford University Press, Oxford (2002) Translated from the French by Reinie Ern e, Oxford Science Publications
95. Lorenzini, D., Tucker, T.J.: Thue equations and the method of Chabauty-Coleman. *Invent. Math.* **148**(1), 47–77 (2002)
96. Luo, Y.: Rank-determining sets of metric graphs. *J. Combin. Theory Ser. A* **118**, 1775–1793 (2011)
97. Luo, Y., Manjunath, M.: Smoothing of limit linear series of rank one on saturated metrized complexes of algebraic curves. Preprint, arXiv:1411.2325 (2014)
98. Maclagan, D., Sturmfels, B.: *Introduction to Tropical Geometry*. Graduate Studies in Mathematics, vol. 161, pp. xii+363. American Mathematical Society, Providence (2015)
99. McCallum, W., Poonen, B.: The method of Chabauty and Coleman. In: *Explicit Methods in Number Theory*. Panorama Synth eses, vol. 36 pp. 99–117. Soci et  Math matique de France, Paris (2012)
100. Merel, L.: Bornes pour la torsion des courbes elliptiques sur les corps de nombres. *Invent. Math.* **124**(1–3), 437–449 (1996)
101. Mikhalkin, G., Zharkov, I.: Tropical curves, their Jacobians and theta functions. In: *Curves and Abelian Varieties*. Contemporary Mathematics, pp. 203–230, vol. 465. American Mathematical Society, Providence, RI (2008)
102. Mumford, D.: An analytic construction of degenerating abelian varieties over complete rings. *Compositio Math.* **24**, 239–272 (1972)
103. Neeman, A.: The distribution of Weierstrass points on a compact Riemann surface. *Ann. Math. (2)* **120**(2), 317–328 (1984)
104. Nicaise, J.: Singular cohomology of the analytic Milnor fiber, and mixed Hodge structure on the nearby cohomology. *J. Algebraic Geom.* **20**(2), 199–237 (2011)
105. Noether, M.: Zur Grundlegung der Theorie der algebraische raumcurven. *J. Reine Angew. Math.* **93**, 271–318 (1882)
106. Ogg, A.P.: On the Weierstrass points of $X_0(N)$. III. *J. Math.* **22**(1), 31–35 (1978)

107. Osserman, B.: Limit linear series for curves not of compact type. Preprint, arXiv:1406.6699 (2014)
108. Payne, S.: Analytification is the limit of all tropicalizations. *Math. Res. Lett.* **16**(3), 543–556 (2009)
109. Perkinson, D., Perlman, J., Wilmes, J.: Primer for the algebraic geometry of sand-piles. In: *Tropical and Non-Archimedean Geometry. Contemporary Mathematics*, vol. 605, pp. 211–256. American Mathematical Society, Providence, RI (2013)
110. Rabinoff, J.: Tropical analytic geometry, Newton polygons, and tropical intersections. *Adv. Math.* **229**(6), 3192–3255 (2012)
111. Raynaud, M.: Spécialisation du foncteur de Picard. *Inst. Hautes Études Sci. Publ. Math.* **(38)**, 27–76 (1970)
112. Severi, F.: Sulla classificazione delle curve algebriche e sul teorema di esistenza di Riemann. *Rend. R. Acc. Naz. Lincei* **24**(5), 887–888 (1915)
113. Stoll, M.: Independence of rational points on twists of a given curve. *Compos. Math.* **142**(5), 1201–1214 (2006)
114. Stoll, M.: Uniform bounds for the number of rational points on hyperelliptic curves of small Mordell-Weil rank. *J. Eur. Math. Soc.* arXiv:1307.1773v4 (2013)
115. Teixidor i Bigas, M.: Injectivity of the symmetric map for line bundles. *Manuscripta Math.* **112**(4), 511–517 (2003)
116. Thuillier, A.: Théorie du potentiel sur les courbes en géométrie analytique non archimédienne. Applications à la théorie d’Arakelov. Ph.D. thesis, University of Rennes (2005)
117. Thuillier, A.: Géométrie toroïdale et géométrie analytique non archimédienne. *Manuscripta Math.* **123**(4), 381–451 (2007)
118. Viviani, F.: Tropicalizing vs. compactifying the Torelli morphism. In: *Tropical and Non-Archimedean Geometry. Contemporary Mathematics*, vol. 605, pp. 181–210. American Mathematical Society, Providence, RI (2013)
119. Voisin, C.: Sur l’application de Wahl des courbes satisfaisant la condition de Brill-Noether-Petri. *Acta Math.* **168**(3–4), 249–272 (1992)
120. Zhang, S.: Admissible pairing on a curve. *Invent. Math.* **112**(1), 171–193 (1993)

Matroid Theory for Algebraic Geometers

Eric Katz

Abstract This article is an introduction to matroid theory aimed at algebraic geometers. Matroids are combinatorial abstractions of linear subspaces and hyperplane arrangements. Not all matroids come from linear subspaces; those that do are said to be *representable*. Still, one may apply linear algebraic constructions to non-representable matroids. There are a number of different definitions of matroids, a phenomenon known as *cryptomorphism*. In this survey, we begin by reviewing the classical definitions of matroids, develop operations in matroid theory, summarize some results in representability, and construct polynomial invariants of matroids. Afterwards, we focus on matroid polytopes, introduced by Gelfand–Goresky–MacPherson–Serganova, which give a cryptomorphic definition of matroids. We explain certain locally closed subsets of the Grassmannian, thin Schubert cells, which are labeled by matroids, and which have applications to representability, moduli problems, and invariants of matroids following Fink–Speyer. We explain how matroids can be thought of as cohomology classes in a particular toric variety, the *permutohedral variety*, by means of Bergman fans, and apply this description to give an exposition of the proof of log-concavity of the characteristic polynomial of representable matroids due to the author with Huh.

Keywords Matroid theory • Intersection theory • Toric varieties

1 Introduction

This survey is an introduction to matroids for algebraic geometry-minded readers. Matroids are a combinatorial abstraction of linear subspaces of a vector space with distinguished basis or, equivalently, a set of labeled set of vectors in a vector space. Alternatively, they are a generalization of graphs and are therefore amenable to a structure theory similar to that of graphs. For algebraic geometers, they are a

E. Katz (✉)

Department of Combinatorics & Optimization, University of Waterloo, Waterloo, ON,
Canada N2L 3G1

e-mail: eeekatz@math.uwaterloo.ca

source of bizarre counterexamples in studying moduli spaces, a combinatorial way of labelling strata of a Grassmannian, and a testing ground for theorems about representability of cohomology classes.

Matroids were introduced by Whitney [108] as an abstraction of linear independence. If \mathbf{k} is a field, one can study an $(n + 1)$ -tuple of vectors (v_0, \dots, v_n) of \mathbf{k}^{d+1} by defining a *rank function*

$$r : 2^{\{0, \dots, n\}} \rightarrow \mathbb{Z}_{\geq 0}$$

by, for $S \subset \{0, \dots, n\}$,

$$r(S) = \dim(\text{Span}(\{v_i \mid i \in S\})).$$

This rank function satisfies certain natural properties, and one can consider rank functions that satisfy these same properties without necessarily coming from a set of vectors. This rank function is what Whitney called a matroid. Whitney noticed that there were matroids that did not come from a set of vectors over a particular field \mathbf{k} . Such matroids are said to be *non-representable* over \mathbf{k} . Matroids can also be obtained from graphs as a sort of combinatorial abstraction of the cycle space of a graph. In fact, it is a piece of folk wisdom that any theorem about graph theory that makes no reference to vertices is a theorem in matroid theory. The important structure theory of matroids that are representable over particular finite fields (or over all fields) was initiated by Tutte. In fact, Tutte was able to demarcate the difference between graphs and matroids in a precise way. The enumerative theory of matroids as partially ordered sets was initiated by Birkhoff, Whitney, and Tutte and systematized and elaborated by Rota [82]. From this enumerative theory, there were associated polynomial invariants of matroids, among them the characteristic and Tutte polynomials. The theory of matroids was enlarged and formulated in more categorical terms by a number of researchers including Brylawski, Crapo, Higgs, and Rota [22]. Matroids were found to have applications to combinatorial optimization problems by Edmonds [32] who introduced a polytope encoding the structure of the matroid.

Throughout the development of the subject, many alternative formulations of matroids were found. They were combinatorial abstractions of notions like the span of a subset of a set of vectors, independent sets of vectors, vectors forming a basis of the ambient space, or minimally dependent sets of vectors. Each of these definitions made different structures of matroids more apparent. The multitude of non-obviously equivalent definitions goes by the name of *cryptomorphism*. Matroid theory is, in fact, sometimes forbidding to beginners because of the frequent switching between definitions.

The point of view of matroids in this survey is one initiated in the work of Gelfand–Goresky–MacPherson–Serganova, which relates representable matroids to certain subvarieties of a Grassmannian. One views the vector configuration spanning \mathbf{k}^{d+1} as a surjective linear map

$$\mathbf{k}^{n+1} \rightarrow \mathbf{k}^{d+1}, (x_0, \dots, x_n) \mapsto x_0 v_0 + \dots + x_n v_n.$$

By dualizing, one has an injective linear map $(\mathbf{k}^{d+1})^* \rightarrow (\mathbf{k}^{n+1})^*$ where we consider the image of the map as a subspace $V \subseteq (\mathbf{k}^{n+1})^* \cong \mathbf{k}^{n+1}$. Now, we may scale this subspace by the action of the algebraic torus $(\mathbf{k}^*)^{n+1}$ acting on the coordinates of \mathbf{k}^{n+1} . The closure of the algebraic torus orbit containing V defines a subvariety of $\text{Gr}(d + 1, n + 1)$. The set of characters of the algebraic torus acting on V leads to the notion of matroid polytopes and a new perspective on matroids. It is this perspective that we will explore in this survey. In particular, we will study how the matroid polytope perspective leads to a class of valutive invariants of matroids, how the subvarieties of $\text{Gr}(d + 1, n + 1)$ that correspond to subspaces V representing a particular matroid are interesting in their own right and give insight to representability, and finally how the study of an object, called the Bergman fan, parameterizing degenerations of the matroid sheds light on the enumerative theory of matroids. In fact, we will use the Bergman fan to present a proof of a theorem of Huh and the author [47] of a certain set of inequalities among coefficients of the characteristic polynomial of matroids, called log-concavity, addressing part of a conjecture of Rota, Heron, and Welsh [83].

Another theme of this survey is cryptomorphism. The new ways of thinking about matroids introduced by algebraic geometry have introduced two new definitions of matroids that are quite different from the more classical characterizations: matroid polytopes and Bergman fans. The definition of matroids in terms of matroid polytopes comes out of the work of Gelfand–Goresky–MacPherson–Serganova. The Bergman fan definition was motivated by valuation theory [8], rephrased in terms of tropical geometry, and can be described as studying matroids as cohomology classes on a particular toric variety called the permutohedral variety. In this survey, we advocate for the combinatorial study of a slight enlargement of the category of matroids, that of Minkowski weights on permutohedral varieties. This enlargement allows a new operation called (r_1, r_2) -truncation introduced by Huh and the author which is essential to the proof of log-concavity of the characteristic polynomial.

We have picked topics to appeal to algebraic geometers. We have put some emphasis on representability, which through Mnëv’s theorem and Vakil’s work on Murphy’s Law plays a central role in constructing pathological examples in algebraic geometry.

This survey’s somewhat bizarre approach and assumptions of background reflect how the author learned the subject. We assume, for example, that the reader is familiar with toric varieties and K -theory but provide an introduction to Möbius inversion. We include just enough of the highlights of the structure theory of matroids to give readers a sense of what is out there. The literature on matroids is vast and the author’s ignorance keeps him from saying more. The more purely combinatorial research in matroid theory has a quite different flavor from our survey. Also, there are a number of topics that would naturally fit into this survey that we had to neglect for lack of expertise. Such topics include Coxeter matroids [15], oriented matroids [14], matroids over a ring as defined by Fink and Moci [35],

hyperplane arrangements [73], and tropical linear subspaces [89]. This survey is rather ahistorical. We neglect nearly all the motivation coming from graph theory.

We make no claims towards originality in this survey. The presentation of Huh–Katz’s proof of log-concavity of the characteristic polynomial differs from that of the published paper [47] but is similar to the exposition in Huh’s thesis [45].

We would like to acknowledge Matthew Baker, Graham Denham, Michael Falk, Alex Fink, June Huh, Sam Payne, Margaret Readdy, Hal Schenck, and Frank Sottile for valuable conversations. This survey arose from an expository talk given at the Simons Symposium on Non-Archimedean and Tropical Geometry. We’d like to thank the organizers of the symposium for their invitation and encouragement.

There are a number of references that we can recommend enthusiastically and which were used extensively in the writing of this survey. Oxley’s textbook [74] is invaluable as a guide to the combinatorial theory. Welsh’s textbook [102] is very broad and geometrically oriented. Wilson [109] gives a nice survey with many examples. Reiner’s lectures [80] explain the theory of matroids and oriented matroids in parallel while also providing historical background. The three Encyclopedia of Mathematics and its Applications volumes, *Combinatorial Geometries* [106], *Theory of Matroids* [105], and *Matroid Applications* [107] are collections of valuable expository articles. In particular, we found [16] very helpful in the writing of this survey. Denham’s survey on hyperplane arrangements [25] is a useful reference for more advanced topics. Thin Schubert cells are treated in detail in [23].

1.1 Notation

We will study algebraic varieties over a field \mathbf{k} . We will use \mathbf{k} to denote $\mathbb{A}_{\mathbf{k}}^1$, the affine line over \mathbf{k} , and we will use \mathbf{k}^* to denote $(\mathbb{G}_m)_{\mathbf{k}}$, the multiplicative group over \mathbf{k} . For a vector space V , V^* will be the dual space. Consequently, $(\mathbf{k}^n)^*$ will be a vector space and $(\mathbf{k}^*)^n$ will be a multiplicative group. We will refer to such $(\mathbf{k}^*)^n$ as an algebraic torus. Our conventions are geared towards working in projective space: matroids will usually be rank $d + 1$ on a ground set $E = \{0, 1, \dots, n\}$.

This survey is organized largely by the mathematical techniques employed. The first six sections are largely combinatorial while the next six are increasingly algebraic. Section 2 provides motivation for the definition of matroids which is given in Sect. 3. Section 4 provides examples while Sect. 5 explains constructions in matroid theory with an emphasis on what the constructions mean for representable matroids viewed as vector configurations, projective subspaces, or hyperplane arrangements. Section 6 discusses representability of matroids. Section 7 introduces polynomial invariants of matroids, in particular the Tutte and characteristic polynomial. Section 8 reviews the matroid polytope construction of Gelfand–Goresky–MacPherson–Serganova and the valuative invariants that it makes possible. Section 9 reviews constructions involving the Grassmannian, describing the relationships between Plücker coordinates and the matroid axioms and between the matroid polytope and torus orbits, and then it discusses realization

spaces and finally, the K -theoretic matroid invariants of Fink and Speyer. Section 10 is a brief interlude reviewing toric varieties. Section 11 introduces Bergman fans and shows that they are Minkowski weights. Section 12 gives a proof of log-concavity of the characteristic polynomial through intersection theory on toric varieties. Section 13 points out some future directions.

2 Matroids as Combinatorial Abstractions

A matroid is a combinatorial object that captures properties of vector configurations or equivalently, hyperplane arrangements. We will informally discuss different ways of thinking about vector configurations as motivation for the rest of the survey. This section is provided solely as motivation and will not introduce any definitions needed for the rest of the paper.

Let \mathbf{k} be a field, and let v_0, v_1, \dots, v_n be vectors in \mathbf{k}^{d+1} that span \mathbf{k}^{d+1} . We can study the dimension of the span of a subset of these vectors. Specifically, for $S \subset \{0, 1, \dots, n\}$, we set

$$V_S = \text{Span}(\{v_i \mid i \in S\}),$$

and we define a *rank function* $r : 2^E \rightarrow \mathbb{Z}$, by for $S \subseteq \{0, 1, \dots, n + 1\}$,

$$r(S) = \dim(V_S).$$

There are some obvious properties that r satisfies: we must have $0 \leq r(S) \leq |S|$; r must be non-decreasing on subsets (so $S \subseteq U$ implies $r(S) \leq r(U)$); and $r(\{0, 1, \dots, n\}) = \dim(\mathbf{k}^{d+1}) = d + 1$. There is a less obvious property: because for $S, U \subseteq \{0, 1, \dots, n\}$, $V_{S \cap U} \subseteq V_S \cap V_U$, we must have

$$r(S \cap U) \leq \dim(V_S \cap V_U) = \dim(V_S) + \dim(V_U) - \dim(V_{S \cup U}) = r(S) + r(U) - r(S \cup U).$$

We can take this one step further and study all rank functions that satisfy these properties. Such a rank function, we will call a *matroid*. Not all matroids will come from vector configurations. Those that do will be said to be *representable*. Remarkably, a number of geometric constructions will work for matroids regardless of their representability.

Instead of studying rank functions, we can study certain collections of subsets of $\{0, 1, \dots, n\}$ that capture the same combinatorial data. We can study *bases* which are $(d + 1)$ -element subsets of $\{0, 1, \dots, n\}$ corresponding to subsets of $\{v_0, v_1, \dots, v_n\}$ that span \mathbf{k}^{d+1} . Or we can study *independent sets* which are subsets of $\{0, 1, \dots, n\}$ corresponding to linearly independent sets of vectors. We can study *circuits* which are minimal linearly dependent sets of vectors. Or we can study *flats*

which correspond to subspaces spanned by some subset $\{v_0, v_1, \dots, v_n\}$. Each of these collections of subsets can be used to give a definition of a matroid.

Alternatively, we can consider linear subspaces instead of vector configurations. Let $V \subseteq \mathbf{k}^{n+1}$ be a $(d + 1)$ -dimensional subspace that is not contained in any coordinate hyperplane. If e_0, e_1, \dots, e_n is the standard basis of \mathbf{k}^{n+1} , then the dual basis $e_0^*, e_1^*, \dots, e_n^*$ induce linear forms on V . This gives a vector configuration in V^* . We can projectivize V to obtain a projective subspace $\mathbb{P}(V) \subset \mathbb{P}^n$. We can rephrase the data of the rank function in terms of a hyperplane arrangement of $\mathbb{P}(V)$. Indeed, if the e_i^* 's are all non-zero, each linear form e_i^* vanishes on a hyperplane, H_i on $\mathbb{P}(V)$, giving a hyperplane arrangement. Note that H_i is the intersection of $\mathbb{P}(V)$ with the coordinate hyperplane in \mathbb{P}^n cut out by $X_i = 0$ where X_i is a homogeneous coordinate. Now, we can define the rank function by, for $S \subset \{0, 1, \dots, n\}$,

$$r(S) = \text{codim} \left(\left(\bigcap_{i \in S} H_i \right) \subset \mathbb{P}(V) \right).$$

There are interpretations of bases, independent sets, circuits, and flats in this language as well.

There are a number of invariants of matroids that correspond to remembering the matroids up to a certain equivalence class, analogous to passing to a Grothendieck ring. One invariant, the Tutte polynomial can be related to the class of the matroid in a Grothendieck ring whose equivalence relation comes from deletion and contraction operations. In the hyperplane arrangement language, deletion corresponds to forgetting a hyperplane on the projective subspace, and contraction corresponds to restricting the arrangement to a hyperplane.

There are geometric constructions that can be performed on the linear subspace V . These constructions can be studied not just for linear subspaces but for matroids. Some of these constructions, we shall see, can be used to give new combinatorial abstractions of linear subspaces and therefore, new definitions of matroids.

One construction involves the Grassmannian. A linear subspace V corresponds to a point in the Grassmannian $\text{Gr}(d + 1, n + 1)$ parameterizing $(d + 1)$ -dimensional subspaces of \mathbf{k}^{n+1} . Certain information about this point is equivalent to the data of the matroid. To speak of it, we have to study the geometry of the Grassmannian. The Grassmannian has a Plücker embedding,

$$i : \text{Gr}(d + 1, n + 1) \hookrightarrow \mathbb{P}^N$$

where $N = \binom{n+1}{d+1} - 1$. The homogeneous coordinates p_B on \mathbb{P}^N are labeled by $B \subset \{0, 1, \dots, n\}$ with $|B| = d + 1$ and are called Plücker coordinates. The data of the matroid is captured by which Plücker coordinates are non-zero. Alternatively, the data can be phrased in terms of a certain group acting on the Grassmannian. Let $T = (\mathbf{k}^*)^{n+1}$ act on \mathbf{k}^{n+1} by dilating the coordinates. This induces an action on the Grassmannian by, for $t \in T$, taking $V \in \text{Gr}(d + 1, n + 1)$ to $t \cdot V = \{tv \mid v \in V\}$.

Note that the diagonal torus of T acts trivially on $\text{Gr}(d + 1, n + 1)$. This group action extends to the ambient \mathbb{P}^N . Given $V \in \text{Gr}(d + 1, n + 1) \subset \mathbb{P}^N$, we can lift V to a point $\tilde{V} \in \mathbf{k}^{N+1}$ and ask what are the characters of T of the smallest subrepresentation of T containing \tilde{V} . This set of characters captures exactly the data of the matroid. Alternatively, we can rephrase this data in terms of group orbits. The closure of the T -orbit containing V , $\overline{T \cdot V} \subset \mathbb{P}^N$ is a polarized projective toric variety. By well-known results in toric geometry, this toric variety corresponds to a polytope. The data of this polytope also corresponds to the matroid. Moreover, one can study polytopes that arise in this fashion combinatorially and even associate them with non-representable matroids. These matroid polytopes can be used to produce an interesting class of invariants of matroids, called valuative invariants. These invariants are those that are well behaved under subdivision of the matroid polytope into smaller matroid polytopes. The K -theory class of the structure sheaf of the closure of the torus orbit, $\mathcal{O}_{\overline{T \cdot V}} \in K_0(\text{Gr}(d + 1, n + 1))$ is a valuative invariant of the matroid introduced by Speyer. By a combinatorial description of K -theory of the Grassmannian, the invariant can be extended to describe non-representable matroids.

Given a matroid M , one can study the set of points on the Grassmannian that have M as their matroid. This locally closed subset of $\text{Gr}(d + 1, n + 1)$ is called a *thin Schubert cell*. Such sets have arbitrarily bad singularities (up to an equivalence relation) and can be used to construct other pathological moduli spaces. The pathological nature of thin Schubert cells are responsible for some difficulties in understanding representability of matroids.

Given a subspace $\mathbb{P}(V) \subset \mathbb{P}^n$, one can blow up the ambient \mathbb{P}^n to understand $\mathbb{P}(V)$ as a homology class. Homogeneous coordinates on \mathbb{P}^n provide a number of distinguished subspaces. In fact, we can consider all subspaces that occur as $\bigcap_{i \in S} H_i$ for all proper subsets $S \subset \{0, 1, \dots, n\}$. If we intersect n distinct coordinate hyperplanes of \mathbb{P}^n , we get a point all but one of whose coordinates are 0. There are $n + 1$ such points. If we intersect $n - 1$ distinct coordinate hyperplanes, we get a coordinate line between two of those points. If we intersect $n - 2$ distinct coordinate hyperplanes, we get a coordinate plane containing three of those points, and so on. We can produce a new variety, called the permutohedral variety, X by first blowing up the $n + 1$ points, then blowing up the proper transforms of the coordinate lines, then blowing up the proper transforms of the coordinate planes, and so on. The proper transform $\widetilde{\mathbb{P}(V)}$ of $\mathbb{P}(V)$ is an iterated blow-up of intersections of hyperplanes from the induced arrangement. The homology class of $\widetilde{\mathbb{P}(V)}$ in X depends only on the matroid of V . It has been studied as the Bergman fan by Sturmfels and Ardila-Klivans. Moreover, it can be defined for non-representable matroids. The characteristic polynomial, which is a certain specialization of the Tutte polynomial, can be phrased as the answer to an intersection theory problem on the Bergman fan. In the representable case, one may apply intersection-theoretic inequalities derived from the Hodge index theorem to prove inequalities between the coefficients of the characteristic polynomial, resolving part of the Rota–Heron–Welsh conjecture.

We will discuss all this and more below.

3 Matroids

There are many definitions of a matroid. The equivalence of these definitions goes by the buzzword of *cryptomorphism*. They are all structures on a finite set E which will be called the ground set.

The rank formulation of matroids is the one that will be most useful to us in the sequel:

Definition 3.1. A matroid on E of rank $d + 1$ is a function

$$r : 2^E \rightarrow \mathbb{Z}$$

satisfying

- (1) $0 \leq r(S) \leq |S|$,
- (2) $S \subseteq U$ implies $r(S) \leq r(U)$,
- (3) $r(S \cup U) + r(S \cap U) \leq r(S) + r(U)$, and
- (4) $r(\{0, \dots, n\}) = d + 1$.

Definition 3.2. Two matroids M_1, M_2 with M_i on E_i with rank function r_i are said to be *isomorphic* if there is a bijection $f : E_1 \rightarrow E_2$ such that for any $S \subseteq E_1, r_2(f(S)) = r_1(S)$.

The definition of matroids makes sense from the point of view of vector configuration in a vector space. Let $E = \{0, 1, \dots, n\}$. Let \mathbf{k} be a field. Consider the vector space \mathbf{k}^{d+1} together with $n + 1$ vectors $v_0, \dots, v_n \in \mathbf{k}^{d+1}$ spanning \mathbf{k}^{d+1} . For $S \subseteq E$, set

$$r(S) = \dim(\text{Span}(\{v_i \mid i \in S\})).$$

Then r is a matroid. If we write

$$V_S = \text{Span}(\{v_i \mid i \in S\}),$$

then item (3.1) is equivalent to

$$\dim(V_{S \cup U}) \leq \dim(V_S \cap V_U).$$

Note that this inequality may be strict because there may be no subset of E exactly spanning $V_S \cap V_U$. Several of the matroid axioms are simple-minded and obvious while one is non-trivial. This is very much in keeping with the flavor of the subject.

Definition 3.3. A matroid is said to be *representable over \mathbf{k}* if it is isomorphic to a matroid arising from a vector configuration in a vector space over \mathbf{k} . A matroid is said to be *representable* if it is representable over some field. A matroid is said to be *regular* if it is representable over every field.

Regular matroids are much studied in combinatorics. They have a characterization due to Tutte (see [74] for details). Given a representable matroid, we may form the matrix whose columns are the coordinates of the vectors in the vector configuration. A matrix is said to be totally unimodular if each square submatrix has determinant 0, 1, or -1 . It is a theorem of Tutte that regular matroids are those representable by totally unimodular matrices with real entries [74].

Representable matroids are an important class of matroids but are not all of them. In fact, it is conjectured that they are asymptotically sparse among matroids. We will discuss non-representable matroids at length in this survey.

Instead of considering a rank function, we may consider instead the set of flats of the matroids.

Definition 3.4. A flat of r is a subset $S \subseteq E$ such that for any $j \in E, j \notin S, r(S \cup \{j\}) > r(S)$.

We think of flats as the linear subspaces of \mathbf{k}^{n+1} spanned by vectors labeled by a subset of E . We may also axiomatize matroids as a set of flats.

Definition 3.5. A matroid is a collection of subsets \mathcal{F} of a set E that satisfy the following conditions

- (1) $E \in \mathcal{F}$,
- (2) if $F_1, F_2 \in \mathcal{F}$, then $F_1 \cap F_2 \in \mathcal{F}$, and
- (3) if $F \in \mathcal{F}$ and $\{F_1, F_2, \dots, F_k\}$ is the set of minimal members of \mathcal{F} properly containing F , then the sets $F_1 \setminus F, F_2 \setminus F, \dots, F_k \setminus F$ partition $E \setminus F$.

Note that axiom (3) implies that for any flat F and $j \notin F$, there is a unique flat F' containing $F \cup \{j\}$ that does not properly contain any flat properly containing F . We can also encode the data of the flats in terms of a closure operation where the closure of a set $S \subseteq E, \text{cl}(S)$ is the intersection of the flats containing S . The set of flats form a lattice which is a poset equipped with operations that abstract intersection and span. The lattice of flats of M is denoted by $L(M)$. We will let $\hat{0}$ be the minimal flat. Given a collection of flats \mathcal{F} , we may recover the rank function of a set $S \subseteq E$ by setting it to be the length of the longest chain of non-trivial flats properly contained in $\text{cl}(S)$.

We can also axiomatize matroids in terms of their bases. A basis for a vector configuration labeled by E is a subset $B \subseteq E$ such that $\{v_i \mid i \in B\}$ is a basis for \mathbf{k}^{d+1} . In terms of the rank function, a basis is a $(d + 1)$ -element set $B \subseteq E$ with $r(B) = d + 1$.

Definition 3.6. A matroid is a collection \mathcal{B} of subsets of E such that

- (1) \mathcal{B} is nonempty and
- (2) If $B_1, B_2 \in \mathcal{B}$ and $i \in B_1 \setminus B_2$, then there is an element j of $B_2 \setminus B_1$ such that $(B_1 \setminus i) \cup \{j\} \in \mathcal{B}$.

The second axiom is called *basis exchange*. It is a classical property of pairs of bases of a vector space and is due to Steinitz. By applying it repeatedly, we may show that all bases have the same number of elements. This is the rank of the matroid. It is straightforward to go from a rank function to a collection of bases and vice versa.

Another axiomatization comes from the set of independent subsets which should be thought of subsets of E labelling linearly independent subsets. These would be subsets $I \subset E$ such that $r(I) = |I|$.

Definition 3.7. A matroid is a collection of subsets \mathcal{I} of E such that

- (1) \mathcal{I} is nonempty,
- (2) Every subset of a member of \mathcal{I} is a member of \mathcal{I} , and
- (3) If X and Y are in \mathcal{I} and $|X| = |Y| + 1$, then there is an element $x \in X \setminus Y$ such that $Y \cup \{x\}$ is in \mathcal{I} .

Definition 3.8. A *loop* of a matroid is an element $i \in E$ with $r(\{i\}) = 0$. A *pair of parallel points* (i, j) of a matroid are elements $i, j \in E$ such that $r(\{i\}) = r(\{j\}) = r(\{i, j\}) = 1$. A matroid is said to be *simple* if it has neither loops nor parallel points.

For vector configurations, a loop corresponds to the zero vector while parallel points correspond to a pair of parallel vectors.

Definition 3.9. A *coloop* of a matroid is an element $i \in E$ that belongs to every basis.

In terms of vector configurations, a coloop corresponds to a vector not in the span of the other vectors.

A *circuit* in a matroid is a minimal subset of E that is not contained in a basis. For a set of vectors, this should be thought of as a subset $C \subseteq E$ such that the vectors labeled by C are linearly dependent but for any $i \in C$, $C \setminus \{i\}$ is linearly independent. Circuits can be axiomatized to give another definition of matroids.

4 Examples

In this section, we explore difference classes of matroids arising in geometry, graph theory, and optimization.

Example 4.1. The uniform matroid $U_{d+1, n+1}$ of rank $d + 1$ on $n + 1$ elements is defined for $E = \{0, 1, \dots, n\}$ by a rank function $r : 2^E \rightarrow \mathbb{Z}_{\geq 0}$ given by

$$r(S) = \min(|S|, d + 1).$$

It corresponds to a vector configuration v_0, v_1, \dots, v_n given by $n + 1$ generically chosen vectors in a $(d + 1)$ -dimensional vector space. Any set of $d + 1$ vectors is a basis. Note that if the field \mathbf{k} does not have enough elements, then the vectors cannot be chosen generically. For example, because \mathbb{F}_2^2 has only three non-zero elements, $U_{2,4}$ does not arise as a vector configuration over \mathbb{F}_2 . Equivalently, it is not representable over \mathbb{F}_2 .

The matroids on a singleton set will be important below. The matroid $U_{0,1}$ is represented by a vector configuration consisting of $0 \in \mathbf{k}$. The single element of the ground set of $U_{0,1}$ is a loop. On the other hand, $U_{1,1}$ is represented by any non-zero vector in \mathbf{k} . The single element of the ground set of $U_{1,1}$ is a coloop.

Example 4.2. Let x_0, x_1, \dots, x_n be points in $\mathbb{P}^d_{\mathbf{k}}$ not contained in any proper projective subspace. Pick a vector $v_i \in \mathbf{k}^{n+1} \setminus \{0\}$ on the line described by x_i . Then v_0, v_1, \dots, v_n gives a vector configuration in \mathbf{k}^{d+1} and therefore a matroid on $\{0, 1, \dots, n\}$ of rank $d + 1$. Specifically, the rank function, for $S \subset E = \{0, 1, \dots, n\}$ is given by

$$r(S) = \dim(\text{Span}(v_i \mid i \in S)).$$

This vector configuration can be thought of as a surjective map:

$$\begin{aligned} \mathbf{k}^{n+1} &\rightarrow \mathbf{k}^{d+1} \\ (t_0, \dots, t_n) &\mapsto t_0 v_0 + \dots + t_n v_n. \end{aligned}$$

Alternatively, we can dualize this map to get an injective map $(\mathbf{k}^{d+1})^* \rightarrow (\mathbf{k}^{n+1})^*$ whose image is a subspace V . Note here that every nonempty set has positive rank so the matroid has no loops. Coloops are elements i such that the minimal projective subspace containing $\{x_0, x_1, \dots, x_n\} \setminus \{x_i\}$ is of positive codimension. Flats correspond to minimal projective subspaces containing some subset of $\{x_0, x_1, \dots, x_n\}$. Bases are $(d + 1)$ -element subsets of $\{x_0, x_1, \dots, x_n\}$ that are not contained in a proper projective subspace of \mathbb{P}^d .

If x_0, x_1, \dots, x_n are contained in a proper projective subspace, we may replace the ambient space by the minimal projective subspace containing x_0, x_1, \dots, x_n in order to define a matroid.

Example 4.3. Let $V \subseteq \mathbf{k}^{n+1}$ be an $(d + 1)$ -dimensional subspace. Let e_0, e_1, \dots, e_n be a basis for \mathbf{k}^{n+1} . The inclusion $i : V \hookrightarrow \mathbf{k}^{n+1}$ induces a surjection $i^* : (\mathbf{k}^{n+1})^* \rightarrow V^*$. The image of the dual basis, $\{i^* e_0^*, i^* e_1^*, \dots, i^* e_n^*\}$ gives a vector configuration in V^* . We can take its matroid.

Let us use the basis to put coordinates (x_0, x_1, \dots, x_n) on \mathbf{k}^{n+1} . We can define subspaces of V as follows: for $I \subseteq E$, set

$$V_I = \{x \in V \mid x_i = 0 \text{ for all } i \in I\}.$$

Then $r(I) = \text{codim}(\dim V_I \subset V)$. A flat in this case is a subset $F \subset E$ such that for any $G \supset F, V_G \subsetneq V_F$. The lattice of flats is exactly the lattice of subsets of V of the form V_S , ordered under reverse inclusion. An element j is a loop if and only if V is contained in the coordinate hyperplane $x_j = 0$. A basis is a $(d + 1)$ -element subset $B \subseteq \{0, 1, \dots, n\}$ such that $V_B = \{0\}$. An element j is a coloop if and only if V contains the basis vector e_j .

If we replace V and \mathbf{k}^{n+1} by their projectivizations, and V is not contained in any coordinate hyperplane, then the definition still makes sense. In this case we consider an r -dimensional projective subspace $\mathbb{P}(V)$ in $\mathbb{P}(V)$. We consider $\mathbb{P}(V_i)$ as an arrangement of hyperplanes on \mathbb{P}^n [73]. In this case, the subsets $\mathbb{P}(V_F)$ as F ranges over all flats correspond to the different possible intersections of hyperplanes (including the empty set).

Example 4.4. Given $V \subset \mathbf{k}^{n+1}$ as above, we may define a matroid by considering the quotient

$$\pi : \mathbf{k}^{n+1} \rightarrow \mathbf{k}^{n+1}/V.$$

For $I \subset E$, let $\mathbf{k}^I \subset \mathbf{k}^{n+1}$ be the subspace given by

$$\mathbf{k}^I = \text{Span}(\{e_i \mid i \in I\}).$$

We set

$$r(I) = \dim(\pi(\mathbf{k}^I)).$$

This is the matroid given by the vector configuration $\{\pi(e_0), \pi(e_1), \dots, \pi(e_n)\}$.

This example is related to the previous one by matroid duality which we will investigate in Sect. 5.3.

Example 4.5. One can draw simple matroids as point configurations. We imagine the points as lying in some projective space, and if the matroid were representable, they would give a vector configuration as in Example 4.2. For example, if we have a rank 3 matroid, we view the points as spanning a projective plane. We specify the rank 2 flats that contain more than two points by drawing lines through the points. These lines together with all lines between pairs of points are exactly the rank 2 flats. Higher rank matroids can be described by point configurations in higher dimensional spaces where we specify the k -flats that contain more points than what





Fig. 1 The Fano and non-Fano matroids

Fig. 2 The Pappus and non-Pappus matroids



is predicted by rank considerations (e.g., for 3-flats which correspond to planes, we would show the planes that are not merely those containing three non-collinear points or a line and a non-incident point).

The Fano and non-Fano matroids denoted by F_7 and F_7^- , respectively, are pictured in Fig. 1. The Fano matroid is a rank 3 matroid consisting of seven points together with seven lines passing through particular triples of points. A line through three of those points is drawn as a circle. The Fano matroid is representable exactly over fields of characteristic 2. In fact, it is the set of all points and lines in the projective plane over $\mathbb{F}_2, \mathbb{P}_{\mathbb{F}_2}^2$. The non-Fano matroid is given by the same configuration but with the center line removed. We see that as meaning that those three points on that line are no longer collinear. This matroid is representable exactly over fields whose characteristic is different from 2.

Example 4.6. The Pappus and non-Pappus matroids are pictured in Fig. 2. The non-Pappus matroid is obtained from the Pappus matroid by mandating that the three points in the middle not be collinear. This is in violation of Pappus’s theorem which is a theorem of projective geometry over a field. Therefore, the non-Pappus matroid is not representable over any field.

Example 4.7. Let G be a graph with $n + 1$ edges labeled by $E = \{0, 1, \dots, n\}$. We can define the graphic matroid $M(G)$ to be the matroid whose set of bases are G ’s spanning forests. The flats of this matroid are the set of edges F such that F contains any edge whose endpoints are connected by a path of edges in F . Note that the loops of this matroid are the loops of the graphs while the coloops are the bridges. In fact, coloops are sometimes called bridges or isthmuses in the literature.

This matroid comes from a vector configuration. Pick a direction for each edge. Let $C_1(G, \mathbf{k})$ be the vector space of simplicial 1-chains on G considered as a 1-dimensional simplicial complex. There is a basis of $C_1(G, \mathbf{k})$ given by the edges e_0, \dots, e_n with their given orientation. Let $\partial : C_1(G, \mathbf{k}) \rightarrow C_0(G, \mathbf{k})$ be the differential. Then $M(G)$ is given by the vector configuration $\partial(e_0), \dots, \partial(e_n)$. Here, the rank of the matroid is

$$\dim(\partial C_1(G, \mathbf{k})) = \dim C_0(G, \mathbf{k}) - \dim H_0(G, \mathbf{k}) = |V(G)| - \kappa(G)$$

where $\kappa(G)$ is the number of connected components of G .

By Whitney’s 2-isomorphism theorem [74, Theorem 5.3.1], one can reconstruct a connected graph G from $M(G)$ up to two moves, vertex cleaving and Whitney

twists. If G is 3-connected, it can be uniquely reconstructed from $M(G)$. In fact, matroids can be considered to be generalizations of graphs. Tutte [98] stated, “If a theorem about graphs can be expressed in terms of edges and circuits only it probably exemplifies a more general theorem about matroids.”

As a special case, consider K_{d+1} , the complete graph on $d + 1$ vertices. Let us denote its vertices by w_0, \dots, w_d . The edges are denoted by e_{ij} for $0 \leq i < j \leq d$. The differential is given by $\partial e_{ij} = w_i - w_j$. The associated subspace as in Example 4.2 is $dC^0(G, \mathbf{k}) \subset C^1(G, \mathbf{k})$. We can put coordinates y_0, \dots, y_d on $C^0(G, \mathbf{k})$ by taking as a basis the characteristic functions of vertices $\delta_{w_0}, \dots, \delta_{w_d}$. Therefore,

$$dC_0(G, \mathbf{k}) \cong \mathbf{k}^{n+1}/\mathbf{k}$$

where we quotient by the diagonal line. The hyperplane arrangement induced by the coordinate subspaces of $C^1(G, \mathbf{k})$ is the *braid arrangement*

$$\{y_i - y_j = 0 \mid 0 \leq i < j \leq d\}.$$

Example 4.8. Let G be a graph with $n + 1$ edges labeled by $E = \{0, 1, \dots, n\}$. We can also define another matroid, the cographic matroid $M^*(G)$ of G . The bases of $M^*(G)$ are complements of the spanning forests of G .

The cographic matroid of a graph also comes from a vector configuration. We pick a direction for each edge as before. Let $C^1(G, \mathbf{k})$ be the 1-cochains of G . It has a basis given by $\delta_{e_0}, \delta_{e_1}, \dots, \delta_{e_n}$, the characteristic function of each edge with given orientation. Let $d : C^0(G, \mathbf{k}) \rightarrow C^1(G, \mathbf{k})$ be the differential. Let $V = C^1(G, \mathbf{k})/dC^0(G, \mathbf{k})$. The matroid is given by the image of $\delta_{e_0}, \delta_{e_1}, \dots, \delta_{e_n}$ in V .

Matroids that are isomorphic to $M(G)$ for some graph G are said to be graphic. Cographic matroids are defined analogously. Observe that because cycle and cocycle spaces can be defined over any field, graphic and cographic matroids are regular. The uniform matroid $U_{2,4}$, because it is not regular, is neither graphic nor cographic. There are examples of regular matroids that are neither graphic nor cographic.

Example 4.9. Transversal matroids arise in combinatorial optimization. Let $E = \{0, \dots, n\}$. Let A_1, \dots, A_m be subsets of E . A *partial transversal* is a subset $I \subseteq E$ such that there exists an injective $\phi : I \rightarrow \{1, \dots, m\}$ such that $i \in A_{\phi(i)}$. A transversal is a partial transversal of size m . We can view elements of E as people and elements of $\{1, \dots, m\}$ as jobs where A_i is the set of people qualified to do job i . A partial transversal is a set of people who can each be assigned to a different job. The transversal matroid is defined to be the matroid on E whose independent sets are the sets of partial transversals.

An important generalization of Hall’s theorem is due to Rado. See [102, Chap. 7] for details:

Theorem 4.10. *Let M be a matroid on a set E . A family of subsets $A_1, \dots, A_m \subseteq E$ has a transversal that is an independent set in M if and only if for all $J \subseteq \{1, \dots, m\}$,*

$$r(\cup_{j \in J} A_j) \geq |J|.$$

This has applications to finding a common transversal for two collections of subsets.

Example 4.11. Algebraic matroids come from field extensions. Let \mathbb{F} be a field and let \mathbb{K} be an extension of \mathbb{F} generated by a finite subset $E \subset \mathbb{K}$. We define a rank function as follows: for $S \subseteq E$,

$$r(S) = \text{tr. deg}(\mathbb{F}(S)/\mathbb{F}),$$

the transcendence degree of $\mathbb{F}(S)$ over \mathbb{F} . It turns out that r defines a matroid. It can be shown that every matroid representable over some field is an algebraic matroid over that same field [74, Proposition 6.7.11]. There are examples of non-algebraic matroids and of algebraic, non-representable matroids.

Algebraic matroids can be interpreted geometrically. Let $V \subset \mathbb{A}_{\mathbb{F}}^{n+1}$ be an algebraic variety. Set $\mathbb{K} = \mathbb{K}(V)$, the function field of V . Let $E = \{x_0, x_1, \dots, x_n\}$ be the coordinate functions on V . The rank of the algebraic matroid given by \mathbb{K} is the transcendence degree of \mathbb{K} over \mathbb{F} , which is equal to $\dim V$. A subset $S \subseteq E$ of size $\dim V$ is a basis when $\text{tr. deg}(\mathbb{K}/\mathbb{F}(S)) = 0$ which happens when the projection onto the coordinate space $\pi_S : V \rightarrow \mathbb{A}_{\mathbb{F}}^S$ is generically finite.

Example 4.12. An important class of matroids are the paving matroids. Conjecturally, they form almost all matroids [68]. A paving matroid of rank $r + 1$ is a matroid such that any set $I \subset E$ with $|I| \leq r$ is independent. Paving matroids can be specified in terms of their hyperplanes, that is, their rank r flats as there is a cryptomorphic axiomatization of matroids in terms of their hyperplanes. We use the following proposition where an r -partition of a set E is a collection of subsets $\mathcal{T} = \{T_1, T_2, \dots, T_k\}$ where for all $i, |T_i| \geq r$ and each r element subset of E is a subset of a unique T_i :

Proposition 4.13 ([74, Proposition 2.1.24]). *If $\mathcal{T} = \{T_1, T_2, \dots, T_k\}$ is an r -partition of a set E , then \mathcal{T} is the set of hyperplanes of a rank $r + 1$ matroid on E . Moreover, for $r \geq 1$, the set of hyperplanes of every rank $r + 1$ paving matroid on E is an r -partition of E .*

One example of a paving matroid is the Vámos matroid. Here we follow the definition of [74]. We set $E = \{1, 2, 3, 4, 1', 2', 3', 4'\}$. We set

$$\mathcal{T}_1 = \{\{1, 2, 1', 2'\}, \{1, 3, 1', 3'\}, \{1, 4, 1', 4'\}, \{2, 3, 2', 3'\}, \{2, 4, 2', 4'\}$$

and

$$\mathcal{T} = \mathcal{T}_1 \cup \{T \subset E \mid |T| = 3 \text{ and } T \text{ is not contained in any element of } \mathcal{T}_1\}$$

Then \mathcal{T} is a 3-partition and the set of hyperplanes of a rank 4 matroid V_8 . This matroid is not representable over any field. One can view V_8 as the set of vertices of a cube whose bottom and top faces are labeled $\{1, 3, 2, 4\}$ and $\{1', 3', 2', 4'\}$ where we mandate that none of the following quadruples of points are coplanar: $\{1, 3, 2, 4\}$, $\{1', 3', 2', 4'\}$, and $\{3, 4, 3', 4'\}$. This matroid turns out to be non-algebraic [49].

For details on paving matroids, see [74, Sect. 2.1].

Example 4.14. Schubert matroids are the matroids whose linear subspaces correspond to the generic point of a particular Schubert cell. They were introduced by Crapo [20]. See [4] for more details.

Schubert cells form an open stratification of the Grassmannian $\text{Gr}(d + 1, n + 1)$ of $(d + 1)$ -dimensional subspaces of \mathbf{k}^{n+1} . The Schubert cells consist of all the subspaces that intersect a flag of subspaces in particular dimensions. Specifically, we have a flag of subspaces

$$\{0\} \subsetneq W_1 \subsetneq W_2 \subsetneq \dots \subsetneq W_n \subsetneq W_{n+1} = \mathbf{k}^{n+1}$$

with $\dim W_i = i$. For a subspace $V \in \text{Gr}(d + 1, n + 1)$, this flag induces a nested sequence of subspaces

$$\{0\} \subseteq V \cap W_1 \subseteq V \cap W_2 \subseteq \dots \subseteq V \cap W_n \subseteq V \cap W_{n+1} = V$$

The dimensions of these subspaces increase by 0 or 1 with each inclusion, so we can mandate where the jump occurs. In the most generic situation, the sequence of dimensions would be $0, \dots, 0, 1, 2, \dots, d, d + 1$, that is, $\dim(V \cap W_{n-d+i}) = i$. We consider cases where the sequence of the jumps differs from that situation. Specifically, we let a_1, \dots, a_{d+1} be a non-increasing sequence of integers with $a_1 \leq n - d$. The Schubert cell is cut out by the open conditions that $\dim(V \cap W_{n-d+i-a_i}) = i$. Its closure, the corresponding Schubert variety is cut out by replacing the equality by “ \geq ”.

To put Schubert cells into a matroid context, we must pick a flag. Let $W_i = \text{Span}(e_0, \dots, e_{i-1})$. In other words, W_i is cut out by the system $x_i = \dots = x_n = 0$. Therefore, we expect to have $\{n - a_{d+1}, n - 1 - a_d, \dots, n - d - a_1\}$ as a basis. However, we need to impose more conditions to specify a matroid. We will suppose that V is generic with respect to the flag apart from these conditions. We declare the bases of the matroid to be exactly the subsets $\{s_0, s_1, \dots, s_d\} \subseteq \{0, \dots, n\}$ such that $s_i \leq (n - i) - a_{d+1-i}$. A point of view that will be taken up later is that matroids allow one to specify points in Schubert cells more precisely.

5 Operations on Matroids

In this section, we survey some of the operations for constructing and relating matroids. Our emphasis is on explaining these constructions in the representable case. For more complete references we recommend [16, 74].

5.1 Deletion and Contraction

Let M be a matroid of rank $d + 1$ on a finite set E . For $X \subset E$, we may define the *deletion* $M \setminus X$. The ground set of $M \setminus X$ is $E \setminus X$ with rank function given as follows: for $S \subseteq E \setminus X$,

$$r_{M \setminus X}(S) = r_M(S).$$

If M is represented by a vector configuration v_0, \dots, v_n , $M \setminus X$ is represented by $\{v_i \mid i \notin X\}$. Similarly, if M is represented by a vector space $V \subseteq \mathbf{k}^{n+1}$, $M \setminus X$ is represented by $\pi(V) \subseteq \mathbf{k}^n$ where $\pi : \mathbf{k}^{n+1} \rightarrow \mathbf{k}^{n+1-|X|}$ is given by projecting out the coordinates corresponding to elements of X . Deletion on graphic matroids corresponds to deleting an edge from the graph.

For $T \subseteq E$, we define the *restriction* by

$$M|_T = M \setminus (E \setminus T).$$

If F is a flat of M , then it is easy to see that the lattice of flats $L(M|_F)$ is the interval $[\hat{0}, F]$ in $L(M)$.

A special case which will be of interest below is the deletion of a single element. Let $i \in E$. If i is not a coloop, then there is a basis B of E not containing i . This is a rank $d + 1$ subset of $M \setminus i$, and so $M \setminus i$ is a matroid of rank $d + 1$. The bases of $M \setminus i$ are the bases of M that do not contain i . By considering the matroid as a linear subspace, we see that $\dim \pi(V) = \dim V$.

There is a dual operation to deletion called *contraction*. Let $X \subset E$. We define the contraction M/X to be the matroid on ground set $E \setminus X$ given by for $S \subseteq E \setminus X$,

$$r(S) = r(S \cup X) - r(X).$$

Consequently, M/X is of rank $n + 1 - r(X)$. If F is a flat of M , then $L(M/F)$ is the interval $[F, E]$ in $L(M)$. If M is represented by a vector configuration, v_0, \dots, v_n in \mathbf{k}^{d+1} , M/X is represented by $\{\pi(v_i) \mid i \notin X\}$ where $\pi : \mathbf{k}^{d+1} \rightarrow \mathbf{k}^{d+1}/W$ is the projection and

$$W = \text{Span}(\{v_i \mid i \in X\}).$$

Likewise if M is represented by a subspace $V \subseteq \mathbf{k}^{n+1}$, M/X is given by $V \cap L_X \subseteq L_X$ where L_X is the coordinate subspace given by

$$L_X = \{x \mid x_i = 0 \text{ if } i \in X\}.$$

If $i \in E$ that is not a loop, then the rank of M/i is d . In this case, if M is represented by a subspace V , then V intersects L_i transversely. If i is a loop, the rank of M/i is $d + 1$. In this case, V is contained in the subspace L_i . If i is not a loop, the bases of M/i are

$$\mathcal{B}(M/i) = \{B \setminus \{i\} \mid i \in B, B \text{ is a basis for } M\}.$$

Definition 5.1. A matroid M' is said to be a *minor* of M if it is obtained by deleting and contracting elements of the ground set of M .

Note that a minor of a representable matroid is representable because we have a geometric interpretation of deletion and contraction on a vector arrangement or linear subspace.

5.2 Direct Sums of Matroids

Given two matroids M_1, M_2 on disjoint sets E_1, E_2 , we may produce a direct sum matroid $M_1 \oplus M_2$ on $E_1 \sqcup E_2$. Specifically, for $S_1 \subseteq E_1, S_2 \subseteq E_2$, we define

$$r(S_1 \sqcup S_2) = r_1(S_1) + r_2(S_2).$$

The bases of M are of the form $B_1 \sqcup B_2$ for $B_1 \in \mathcal{B}_1, B_2 \in \mathcal{B}_2$. The circuits of $M_1 \oplus M_2$ are the circuits of M_1 together with the circuits of M_2 . The lattice of flats obeys

$$L(M_1 \oplus M_2) = L(M_1) \times L(M_2)$$

where the underlying set is the Cartesian product and $(F_1, F_2) \leq (F'_1, F'_2)$ if and only if $F_1 \leq F'_1$ and $F_2 \leq F'_2$. If M_1 is represented by vectors $v_0, \dots, v_m \in \mathbf{k}^{d_1+1}$ and M_2 is represented by vectors $w_0, \dots, w_n \in \mathbf{k}^{d_2+1}$, then $M_1 \oplus M_2$ is represented by

$$(v_0, 0), \dots, (v_m, 0), (0, w_0), \dots (0, w_n) \in \mathbf{k}^{d_1+1} \oplus \mathbf{k}^{d_2+1}.$$

If M_1, M_2 are represented by subspace $V_1 \subseteq \mathbf{k}^{n_1+1}, V_2 \subseteq \mathbf{k}^{n_2+1}$, then $M_1 \oplus M_2$ is represented by

$$V_1 \oplus V_2 \subseteq \mathbf{k}^{n_1+1} \oplus \mathbf{k}^{n_2+1}.$$

For graphic matroids, direct sum corresponds to producing a new graph by identifying a vertex in one graph with a vertex in the other.

Every matroid has a decomposition analogous to the decomposition of a graph into 2-connected components. Here, a graph is said to be 2-connected if it cannot be disconnected by removing one vertex. A decomposition into two 2-connected components involves repeatedly removing each cut vertices and then replacing it with a new vertex in each component.

Definition 5.2. A matroid is *connected* if for every $i, j \in E$, there exists a circuit containing i and j .

We can define connected components of the matroid by saying that two elements i, j are in the same connected component if and only if there exists a circuit containing i and j . This is an equivalence relation. It can be stated in terms of bases in the following form which will be important when we study matroid polytopes: two elements i, j are in the same connected component if and only if there exists bases B_1 and B_2 such that $B_2 = (B_1 \setminus \{i\}) \cup \{j\}$. If T_1, \dots, T_k are the connected components of M , then we have a direct sum decomposition

$$M \cong M|_{T_1} \oplus \cdots \oplus M|_{T_k}.$$

Connected matroids are indecomposable under direct sum, and matroids have a unique direct sum decomposition into connected matroids.

Loops and coloops play a particular role in direct sum decompositions. Because the only circuit in which a loop i occurs is $\{i\}$, $\{i\}$ is a connected component. Similarly, because a coloop j does not occur in any circuit, $\{j\}$ is a connected component. Therefore, loops and coloops may be split off from the matroid as in the following proposition:

Proposition 5.3. Any matroid M can be written as a direct sum

$$M \cong M' \oplus (U_{0,1})^{\oplus l} \oplus (U_{1,1})^{\oplus c}$$

where M' has neither loops nor coloops.

5.3 Duality

Duality is a natural operation on matroids that generalizes duality of planar graphs.

Definition 5.4. The dual of a matroid M on E with rank function r is defined to be the matroid on E with rank function given by

$$r^*(S) = r(E \setminus S) + |S| - r(E).$$

This rank function satisfies the axioms of a matroid. If M is rank $d + 1$ on $E = \{0, 1, \dots, n\}$, then M^* is rank $n - d$. The bases of M^* can be seen to be the set

$$\mathcal{B}^* = \{E \setminus B \mid B \in \mathcal{B}\}.$$

Duality interchanges loops and coloops, commutes with direct sum, and takes deletion to contraction: $(M \setminus i)^* = M^*/i$. As an example, we have $(U_{r+1,n+1})^* = U_{n-r,n+1}$.

The dual of a representable matroid is representable. Let M be represented by a vector configuration v_0, v_1, \dots, v_n spanning \mathbf{k}^{d+1} . This configuration can be thought of as a surjection $p : \mathbf{k}^{n+1} \rightarrow \mathbf{k}^{d+1}$ which fits into an exact sequence as

$$0 \longrightarrow \ker(p) \xrightarrow{j} \mathbf{k}^{n+1} \xrightarrow{p} \mathbf{k}^{d+1} \longrightarrow 0 .$$

We can take duals to get an injection $p^* : (\mathbf{k}^{d+1})^* \hookrightarrow (\mathbf{k}^{n+1})^*$. The dual of M is given by the vector configuration corresponding to the quotient

$$(\mathbf{k}^{n+1})^* \rightarrow (\mathbf{k}^{n+1})^*/p^*((\mathbf{k}^{d+1})^*) = \ker(p)^* .$$

Indeed, let w_0, w_1, \dots, w_n be the dual basis of $(\mathbf{k}^{n+1})^*$. For $S \subseteq \{0, \dots, n\}$, the span of $\{w_i \mid i \in S\}$ induces a linear projection $\pi_S : \mathbf{k}^{n+1} \rightarrow \mathbf{k}^S$. Write $\mathbf{k}^{E \setminus S} \subset \mathbf{k}^{n+1}$ for the kernel of that projection. The rank of S in the vector configuration induced by $\{w_0, \dots, w_n\}$ is

$$\begin{aligned} r^*(S) &= \dim(j^* \pi_S^*((\mathbf{k}^S)^*)) \\ &= \dim(\pi_S(\ker(p))) \\ &= \dim((\ker(p) + \mathbf{k}^{E \setminus S})/\mathbf{k}^{E \setminus S}) \\ &= \dim(\mathbf{k}^{E \setminus S} + \ker(p) / \ker(p)) + \dim(\ker(p)) - \dim(\mathbf{k}^{E \setminus S}) \\ &= \dim(p(\mathbf{k}^{E \setminus S})) + \dim(\ker(p)) - \dim(\mathbf{k}^{E \setminus S}) \\ &= r(E \setminus S) + (n - d) - ((n + 1) - |S|) \\ &= r(E \setminus S) + |S| - r(E). \end{aligned}$$

If a matroid is represented by a subspace $V \subset \mathbf{k}^{n+1}$, its dual is represented by $V^\perp \subset (\mathbf{k}^{n+1})^*$ where $V^\perp = \ker((\mathbf{k}^{n+1})^* \rightarrow V^*)$.

By interpreting of $M(G)$ and $M^*(G)$ as vector configurations given by chain and cochain groups, we see that these are dual matroids. If G is a planar graph, it turns out that $M^*(G) = M(G^*)$ where G^* is the planar dual of G . Because one can take the dual of any graphic matroids, the theory of matroids allows one to take the dual of a non-planar graph.

5.4 Extensions

Single-element extension is an operation on matroids inverse to single-element deletion. Its properties were worked out by Crapo. Given a matroid M on a ground set E , it produces a new matroid M' on a ground set $E' = E \sqcup \{p\}$ such that $M = M' \setminus p$. Here we will follow the exposition of [16]. We can partition the flats of M into three sets based on how they change under extension:

$$\mathcal{K}_1 = \{F \in \mathcal{F} \mid F \text{ and } F \cup \{p\} \text{ are both flats of } M'\}$$

$$\mathcal{K}_2 = \{F \in \mathcal{F} \mid F \text{ is a flat of } M' \text{ but } F \cup \{p\} \text{ is not}\}$$

$$\mathcal{K}_3 = \{F \in \mathcal{F} \mid F \cup \{p\} \text{ is a flat of } M' \text{ but } F \text{ is not}\}.$$

This partition can be understood in the case where M' is given by a vector configuration $\{v_0, \dots, v_n, v_p\} \subseteq \mathbf{k}^{d+1}$ and $M = M' \setminus p$. The flats in \mathcal{K}_3 correspond to subspaces that contain v_p and are spanned by a subset of $\{v_0, \dots, v_n\}$. A flat F of M is in \mathcal{K}_1 when $v_p \notin \text{Span}(F)$ and $\text{Span}(F \cup v_p)$ is not among the linear subspaces corresponding to flats of M . Finally, $F \in \mathcal{K}_2$ when $v_p \notin \text{Span}(F)$ but $\text{Span}(F \cup v_p)$ corresponds to some flat of M . This flat must be an element of \mathcal{K}_3 . If one knows \mathcal{K}_3 , one can determine \mathcal{K}_2 and therefore \mathcal{K}_1 : elements of \mathcal{K}_2 are exactly those that are contained in an element of \mathcal{K}_3 . Now, let us determine what properties that \mathcal{K}_1 and \mathcal{K}_3 should have. If $F_1 \subseteq F_2$ and $F_2 \in \mathcal{K}_1$, then $F_1 \in \mathcal{K}_1$. Also, if $F_1 \subseteq F_2$ and $F_1 \in \mathcal{K}_3$, then $F_2 \in \mathcal{K}_3$. A more subtle property can be seen by considering intersections: if $F_1, F_2 \in \mathcal{K}_3$, then $v_p \in \text{Span}(F_1) \cap \text{Span}(F_2)$; therefore, if $\text{Span}(F_1 \cap F_2) = \text{Span}(F_1) \cap \text{Span}(F_2)$, then $F_1 \cap F_2 \in \mathcal{K}_3$. These properties can all be established in the abstract combinatorial setting. Translated into the matroid axioms, these properties say that \mathcal{K}_3 is a modular cut:

Definition 5.5. Let M be a matroid. A subset $\mathcal{M} \subseteq \mathcal{F}$ is said to be a *modular cut* if

- (1) If $F_1 \in \mathcal{M}$ and $F_2 \in \mathcal{F}$ with $F_1 \subseteq F_2$, then $F_2 \in \mathcal{M}$ and
- (2) if $F_1, F_2 \in \mathcal{M}$ satisfy

$$r(F_1) + r(F_2) = r(F_1 \cap F_2) + r(F_1 \cup F_2)$$

then $F_1 \cap F_2 \in \mathcal{M}$.

Note that the condition on ranks in (5.5) above says that $\text{Span}(F_1 \cap F_2) = \text{Span}(F_1) \cap \text{Span}(F_2)$.

This characterization of \mathcal{K}_3 holds for all extensions of matroids and is in fact a sufficient condition for an extension to exist:

Proposition 5.6. For any single-element extension M' of M , the set \mathcal{K}_3 is a modular cut. Moreover, given any modular cut \mathcal{M} of a matroid M , there is a unique single-element extension M' , denoted by $M +_{\mathcal{M}} p$ such that $\mathcal{K}_3 = \mathcal{M}$.

An extension is a composition of single-element extensions. Single-element extensions may leave the class of representable matroids: the matroid M may be representable but M' may not be. Indeed, given a matroid M , there may not exist a set of vectors $\{v_0, \dots, v_n\} \subseteq \mathbf{k}^{d+1}$ representing M such that there is a vector v_p contained in the subspaces corresponding to flats in the modular cut. In fact, one can produce such an M by beginning with a non-representable matroid and deleting elements until it becomes representable.

Single-element extensions have a geometric interpretation when one considers matroids represented by a subspace $V \subseteq \mathbf{k}^{n+1}$. Specifically, one looks for a subspace $V' \subset \mathbf{k}^{n+2}$ such that $\pi(V') = V$ where $\pi : \mathbf{k}^{n+2} \rightarrow \mathbf{k}^{n+1}$ is projection onto the first $n + 1$ factors. We can interpret such a V' as the graph of a linear function $l : V \rightarrow \mathbf{k}$. Then \mathcal{K}_3 can be interpreted as the flats contained in $l^{-1}(0)$.

A special case of single-element extension is that of a *principal extension*. Specifically, one takes the modular cut to be all flats containing a given flat F . In terms of vector configuration, this corresponds to adjoining a generic vector in the subspace corresponding to F . The extension is denoted by $M +_F p$. In the case where $F = E$, this is called the *free extension* and it corresponds to extending by a generic vector.

There is an operation inverse to contraction, called coextension. In other words, given a matroid M , one produces M' such that $M = M'/p$. Note that the rank of M' will be one greater than the rank of M . Because deletion is dual to contraction, coextension can be defined in terms of extension. Specifically, let \mathcal{M} be a modular cut in M^* . Then the coextension associated with \mathcal{M} is

$$M' = (M^* +_{\mathcal{M}} p)^*.$$

One can similarly define principal and free coextension.

5.5 Quotients and Lifts

There is a natural quotient operation on matroids. Specifically, one extends a matroid by an element, and then one contracts the element. The quotient given by a modular cut \mathcal{M} is defined to be $M +_{\mathcal{M}} p/p$. In the case of vector configuration, this corresponds to taking a quotient of the ambient subspace: if $M' = M +_{\mathcal{M}} p$ is represented by $\{v_0, \dots, v_n, v_p\}$ in \mathbf{k}^{n+1} , then M'/p is represented by the image of $\{v_0, \dots, v_n\}$ in $\mathbf{k}^{n+1}/\mathbf{k}v_p$. One can define principal quotients by taking modular cuts associated with a flat. If $\mathcal{M} = \{E\}$, then the quotient is given by a free extension by v_p followed by a contraction by v_p . This is called the *truncation* $\text{Trunc}^d(M)$ of M . It corresponds to taking the quotient of the ambient space by a generic vector. In general, for $0 \leq k \leq d + 1$, we define the k -truncation as the matroid on E with rank function given by

$$r_{\text{Trunc}^k(M)}(S) = \max(r(S), k).$$

If we view the matroid as a linear subspace $V \subset \mathbf{k}^{n+1}$, then truncation has a geometric interpretation as follows:

Lemma 5.7. *Let M be a rank $d + 1$ matroid on $\{0, 1, \dots, n\}$ represented by a $(d + 1)$ -dimensional subspace $V \subseteq \mathbf{k}^{n+1}$. Let H be a hyperplane in \mathbf{k}^{n+1} that intersects V_I transversely for all $I \subset \{0, 1, \dots, n\}$. Then $V \cap H$ represents $\text{Trunc}^d(M)$.*

Proof. We see that the rank function associated with $V \cap H$ obeys

$$r_{V \cap H}(S) = \text{codim}(\dim(V_S \cap H) \subset (V \cap H)) = \max(\text{codim}(V_S \subset V), d) = r_{\text{Trunc}^d(M)}(S).$$

□

Consequently, the k -truncation corresponds to intersecting with a generic $n + 1 - (d + 1 - k)$ -dimensional subspace. This will be important in the sequel. There is also a dual notion to quotient, that of *lifts*.

5.6 Maps

There are several rival notions of morphisms between matroids. We briefly review the notion of *strong maps* following Kung [59] noting that there is also a notion of weak maps.

Definition 5.8. Let M_1 and M_2 be matroids on sets E_1 and E_2 , respectively. Let the direct sums $M_i \oplus U_{0,1}$ be given by extending by a loop o_i . A strong map σ from M_1 to M_2 is a function $\sigma : E_1 \cup \{o_1\} \rightarrow E_2 \cup \{o_2\}$ taking o_1 to o_2 such that the preimage of any flat of $M_2 \oplus U_{0,1}$ is a flat in $M_1 \oplus U_{0,1}$.

A strong map turns out to be the composition of an extension and a contraction.

Definition 5.9. An *embedding* of a matroid M on E into a matroid M^+ on E^+ is an inclusion $i : E \hookrightarrow E^+$ such that $M^+|_{i(E)} = M$ where we identify elements of E with their images under i .

Embeddings are strong maps. Contractions are strong maps as well. If we contract by a set $U \subseteq E$, then the map is given by $\sigma : E \cup \{o_1\} \rightarrow (E \setminus U) \cup \{o_2\}$ where σ is the identity on $E \setminus U$ and takes $U \cup \{o_1\}$ to o_2 . We see that we need to add a loop to define strong maps because we will need a zero element as a target for elements. We have the following factorization theorem:

Theorem 5.10. *Let $\sigma : M_1 \rightarrow M_2$ be a strong map. Then there is a matroid M^+ such that σ can be factored as the composition*

$$\sigma : M_1 \xrightarrow{i} M^+ \xrightarrow{p} M_2$$

where i is an embedding and p is a contraction.

5.7 Relaxation

New matroids can be obtained from old by the technique of relaxation. Recall that a circuit of a matroid is a minimal dependent set. Matroids can be axiomatized in terms of circuits. Let $X \subseteq E$ be a circuit that is also a hyperplane. We call such a set a *circuit-hyperplane*. We construct a new matroid by mandating that X be a basis.

Proposition 5.11 ([74, Proposition 1.5.14]). *Let X be a circuit-hyperplane of a matroid M on E . Let \mathcal{B} be the set of bases of M and set $\mathcal{B}' = \mathcal{B} \cup \{X\}$. Then \mathcal{B}' is the set of bases for a matroid M'*

By relaxing a representable matroid, one may obtain a non-representable matroid. In fact, the non-Pappus matroid is obtained from the Pappus matroid by relaxing one line through three points. The fields over which a matroid is representable may change by relaxation as the non-Fano matroid is a relaxation of the Fano matroid.

6 Representability and Excluded Minor Characterizations

6.1 Introduction to Representability

Representability is a central part of the combinatorial study of matroids. Here, one would like a combinatorial description of representable matroids. This is probably too much to ask. However, one can study matroids that are representable over a fixed field. Here, we need to discuss which classes of matroids might have a good structure theory.

The analogy with graph theory is particularly strong here. The prototypical structural result that one would like to generalize is Wagner's characterization of planar graphs which states that a graph is planar if and only if it does not contain a $K_{3,3}$ or K_5 minor. Note that the class of planar graphs is minor closed, that is, any minor of a graph in this class is also in this class. Wagner's theorem produces a finite list of forbidden minors for this class of graphs. A far-reaching generalization of this theorem is the Robertson–Seymour graph minors theorem which states that any minor-closed class of graphs has a finite list of excluded minors. Examples of minor-closed classes of graphs include trees, linklessly embeddable graphs, and graphs with embedding genus at most g for a fixed g . The theorem gives a structural decomposition of an arbitrary minor-closed class of graphs and takes up more than 500 journal pages. While it shows that the number of forbidden minors is finite, it does not construct an explicit list. See [28] for a nice summary.

Just as one might study particular classes of graphs like planar graphs, one can study particular classes of matroids. The most natural classes of matroids would be those that are closed under taking minors, i.e. deleting and contracting elements. One would hope for forbidden minor characterizations of these classes.

There are a number of natural minor-closed classes of matroids. For a fixed field \mathbf{k} , the class of matroids that are representable over \mathbf{k} is minor-closed. This follows from the explicit construction of deletion and contraction on matroids represented by, say, a vector configuration. Other minor-closed classes of graphs are graphic matroids (those of the form $M(G)$ for some graph G), cographic matroids ($M^*(G)$) which are minor-closed because deletion and contraction correspond to operations on the graphs. Also, regular matroids, that is, matroids that are represented over every field, form a minor closed class.

There are all sorts of interesting containments between these classes of matroids. All graphic matroids are regular, but the converse is not true. For example, $M^*(K_5)$, being cographic, is regular, but it is not graphic while by duality $M(K_5)$ is not cographic. The direct sum, $M(K_5) \oplus M^*(K_5)$ is regular but is neither graphic nor cographic since it contains $M(K_5)$ and $M^*(K_5)$ as minors. Not all matroids are representable over \mathbb{F}_2 . For example, the uniform matroid $U_{2,4}$ is not representable over \mathbb{F}_2 because there do not exist four vectors in general position in \mathbb{F}_2^2 . Moreover, not all \mathbb{F}_2 -representable matroids are regular. For example, the Fano matroid, F_7 is representable over \mathbb{F}_2 without being regular. Indeed, it is representable over fields \mathbf{k} of characteristic 2. The non-Fano matroid, on the other hand, turns out to be representable only over fields of characteristic not equal to 2.

There are well-known forbidden minor characterizations of certain classes of matroids, a line of inquiry initiated by Tutte [97]. Matroids representable over \mathbb{F}_2 are characterized by not having a $U_{2,4}$ -minor by a theorem of Tutte. Graphic matroids are exactly the matroids that do not have a $U_{2,4}$, $M^*(K_{3,3})$, or $M^*(K_5)$ minor. This fact is closely related to Wagner's theorem: because the dual of a graphic matroid is the graphic matroid of the dual graph, if it exists, it turns out that $M^*(K_{3,3})$ and $M^*(K_5)$ are not graphic. One can even phrase Wagner's theorem in the language of matroids: a matroid is the graphic matroid of a planar graph if and only if it does not contain any of the following matroids as minors: $U_{2,4}$, $M(K_{3,3})$, $M^*(K_{3,3})$, $M(K_5)$, $M^*(K_5)$. Matroids representable over \mathbb{F}_3 are those without minors isomorphic to $U_{2,5}$, $U_{3,5}$, F_7 , or F_7^* by a theorem due independently to Bixby [12] and Seymour [85]. By [97], regular matroids are exactly the matroids not containing $U_{2,4}$, F_7 , or F_7^* .

These forbidden minor theorems can be proved by looking at possible matrices whose columns are the vectors of a representation. One figures out the pattern of zero and non-zero entries in the matrix by considering the *fundamental circuits* of the matroid. Specifically, one fixes a basis B , say $B = \{0, 1, \dots, d\}$ and supposes that in the representation, the elements of the basis are given by the standard basis vectors. Then given $i \notin B$, one can look at the circuit containing i and some elements of B . The vector representing i must be a linear combination of the standard basis vectors in the circuit. One considers different possibilities for the non-zero entries. If one begins with a matroid that is minor-minimal among non-representable matroids, then the matrix turns out not to represent the matroid. By applying matrix operations organized by combinatorial operations on the matroid, one can put the matrix and matroid in a standard form. For example, because \mathbb{F}_2 has only one non-zero element, if a matroid is \mathbb{F}_2 -representable, once one fixes a basis, only one matrix needs to be

considered. Many of these representability theorems were originally proved using Tutte's Homotopy Theorem [97], a deep and difficult theorem about the structure of a particular polyhedral complex on the flats of low corank. Current research in representability relies on quite difficult theorems on decomposing matroids. See [74, Chap. 14] for a survey.

These excluded minor characterizations are over a fixed finite field, not over a field of fixed characteristic. In particular, one is not allowed to take algebraic extensions of a given field. This in essence creates two ways in which a matroid may fail to be representable: that there are not enough elements; or there is a contradiction in the linear conditions. For example, the uniform matroid $U_{2,4}$ is not representable over \mathbb{F}_2 because there are no generic planes in \mathbb{F}_2^4 . The second problem is much more serious as can be seen in the matroids built by applying Mněv's theorem as we will discuss in Sect. 9.6.

The question of representability over infinite fields is much more difficult and is connected to Hilbert's Tenth problem, the question of algorithmic decidability of a system of polynomial equations. We will discuss this connection below. It turns out that there are infinitely many excluded minors for real representability. The hopelessness of the situation can be distilled into the following theorem of Mayhew, Newman, and Whittle [67] which gives an explicit construction:

Theorem 6.1. *Let \mathbb{K} be an infinite field. Let N be a matroid that is representable over \mathbb{K} . Then there exists a matroid M that is an excluded minor for \mathbb{K} -representability, is not representable over any field, and has N as a minor.*

There has been research on the question of whether it is possible at all to give a finite, combinatorial condition for a matroid to be representable. This is the so-called missing axiom of matroid theory. A paper of Vámos [100] from 1978 asserted that the "missing axiom of matroid theory is lost forever." In other words, there is no way of axiomatizing representability in second-order logic. More recently, however, a paper of Mayhew, Newman, and Whittle [69] criticized Vámos's paper by stating that the matroids in it differed from matroids as commonly studied and asked again "Is the missing axiom of matroid theory lost forever?" conjecturing that representability cannot be axiomatized in a specific second-order logical language. They answered that question with "yes, the missing axiom of matroid theory is lost forever" in a very recent preprint [70].

The general situation of representability over finite fields is the subject of Rota's conjecture that given any finite field \mathbb{F}_q , there are finitely many forbidden minors for \mathbb{F}_q -representability [83]. A solution has been announced by Geelen, Gerards, and Whittle [41]. In addition, they have announced an excluded minor theorem for classes of \mathbb{F}_q -representable matroids:

Theorem 6.2. *Any minor-closed class of \mathbb{F}_q -representable matroids has a finite set of excluded \mathbb{F}_q -representable minors.*

Their work relies on a deep structure theory of matroids which generalizes the work of Robertson and Seymour. See the survey [40] for an earlier account of their progress on Rota's conjecture.

6.2 Inequivalent Representations

A matroid may have several representations that are essentially different. Here, we follow the exposition of Oxley [74]. Suppose a matroid is represented by two vector configurations $\{v_0, v_1, \dots, v_n\}$ and $\{v'_0, v'_1, \dots, v'_n\}$ in \mathbf{k}^{n+1} . The two representations are said to be inequivalent if there does not exist an element of $\text{GL}_{n+1}(\mathbf{k})$ taking one vector configuration to the other. We may also consider representations of simple matroids as point configurations in \mathbb{P}^n . Then we say two representations are projectively inequivalent if they are not related by an element of $\text{PGL}_{n+1}(\mathbf{k})$.

It is not too hard to find matroids with inequivalent representations. For example, $U_{3,5}$ has projectively inequivalent representations over \mathbb{F}_5 . Perhaps surprisingly, there are cases where matroids have projectively unique representations. There is the following theorem of White:

Theorem 6.3. *Let M be a simple matroid that is representable over \mathbb{F}_2 . Then for any field \mathbf{k} , any two representations of M over \mathbf{k} are projectively equivalent.*

Similarly, matroids have at most one representations over \mathbb{F}_3 up to projective equivalence by a theorem of Brylawski and Lucas. There is much research in trying to bound the number of inequivalent representations of certain classes of matroids.

The existence of inequivalent representations is important for questions of representability: given a matroid M that we wish to represent, we may delete different elements, try to find representations of various deletions, and then glue the representations of the deletions together; the existence of inequivalent representations is an obstruction to gluing the smaller representations together. This phenomenon seems, to this author at least, reminiscent of the issues with automorphisms in moduli theory.

6.3 Ingleton's Criterion

By linear algebraic arguments, Ingleton provided a necessary condition for a matroid to be representable over some field [48].

Theorem 6.4. *Let M be a representable matroid on E . Then for any subsets $X_1, X_2, X_3, X_4 \subseteq E$, the rank function obeys the following inequality:*

$$\begin{aligned} & r(X_1) + r(X_2) + r(X_1 \cup X_2 \cup X_3) + r(X_1 \cup X_2 \cup X_4) + r(X_3 \cup X_4) \\ & \leq r(X_1 \cup X_2) + r(X_1 \cup X_3) + r(X_1 \cup X_4) + r(X_2 \cup X_3) + r(X_2 \cup X_4). \end{aligned}$$

More recently, Kinser discovered an infinite family of independent inequalities whose leading member is Ingleton's [55].

Theorem 6.5. *Let $k \geq 4$ and let M be a representable matroid on E . Let $X_1, X_2, \dots, X_k \subseteq E$ be subsets. Then the rank function obeys the following:*

$$\begin{aligned}
& r(X_1 \cup X_2) + r(X_1 \cup X_3 \cup X_k) + r(X_3) + \sum_{i=4}^k (r(X_i) + r(X_2 \cup X_{i-1} \cup X_i)) \\
& \leq r(X_1 \cup X_3) + r(X_1 \cup X_n) + r(X_2 \cup X_3) + \sum_{i=4}^k (r(X_2 \cup X_i) + r(X_{i-1} \cup X_i)).
\end{aligned}$$

The Vamós matroid, V_8 violates Ingleton’s criterion with the choice of $X_i = \{i, i'\}$ and is therefore not representable over any field [74, Exercise 6.1.7]. The Vamós matroid turns out to be a non-representable matroid on the minimum number of elements.

6.4 Are Most Matroids Not Representable?

In [17], Brylawski and Kelly claimed

“It is an exercise in random matroids to show that most matroids are not coordinatizable [representable] over any field (or even any division ring).”

Unfortunately, to this date no one has been able to complete this exercise. However, one can study matroid properties asymptotically: one examines the proportion of matroids on an n element ground set that have a given property and take the limit as n goes to infinity. It is suspected that asymptotically most matroids are non-representable paving matroids. A good reference for what is known is [68]. It is mentioned there that by counting arguments, for a fixed finite field \mathbb{F} , asymptotically, most matroids are not representable over \mathbb{F} .

7 Polynomial Invariants of Matroids

7.1 The Characteristic and Tutte Polynomials

The characteristic polynomial is an invariant of matroids that generalizes the chromatic polynomial of graphs as introduced by Birkhoff [11]. Let G be a graph. The chromatic polynomial of G is the function χ_G given by setting $\chi_G(q)$ to be the number of proper colorings of G with q colors for $q \in \mathbb{Z}_{\geq 1}$. A proper coloring with q colors is a function $c : V(G) \rightarrow \{1, 2, \dots, q\}$ such that adjacent vertices are assigned different values. The function $\chi_G(q)$ can easily be shown to be a polynomial. Note that a graph with a loop has no proper colorings. Recall that for $e = v_1 v_2$, an edge of G , the deletion $G \setminus e$ is the graph G with the edge e removed while the contraction G/e is the graph G with e, v_1, v_2 contracted to a single vertex. The chromatic polynomial obeys the deletion/contraction relation

$$\chi_G(q) = \chi_{G \setminus e}(q) - \chi_{G/e}(q).$$

This can be seen by interpreting $\chi_{G \setminus e}(q)$ as the count of colorings where we do not impose the condition that the colors of the end-points of e are different and interpreting $\chi_{G/e}(q)$ as colorings where the colors of the end-points of e are mandated to be the same. One can show that $\chi_G(q)$ is a polynomial by applying deletion/contraction repeatedly and then noting that the chromatic polynomial is multiplicative over disjoint union of graphs and that if G consists of a single vertex, then $\chi_G(q) = q$.

The chromatic polynomial can be extended to matroids in a straightforward fashion. One looks for a polynomial $\chi_M(q)$, called *the characteristic polynomial* for a matroid M which satisfies

1. (loop property) if M has a loop, then $\chi_M(q) = 0$,
2. (normalization) the characteristic polynomial of the uniform matroid $U_{1,1}$ satisfies

$$\chi_{U_{1,1}}(q) = q - 1.$$

3. (direct sum) If $M = M_1 \oplus M_2$ then

$$\chi_M(q) = \chi_{M_1}(q)\chi_{M_2}(q),$$

4. (deletion/contraction) if i is not a loop or coloop of M then

$$\chi_M(q) = \chi_{M \setminus i}(q) - \chi_{M/i}(q),$$

By inducting on the size of the ground set, we easily see that an invariant satisfying these properties is unique. Note that loops and coloops split off from a matroid in a direct sum decomposition. It turns out that $\chi_M(q)$ is well defined. In fact, we will construct it explicitly below by using Möbius functions.

The characteristic polynomial is a generalization of the chromatic polynomial in the following sense:

Proposition 7.1. *If G is a graph and $M(G)$ is its matroid, then*

$$\chi_G(q) = q^\kappa \cdot \chi_{M(G)}(q)$$

where κ is the number of components of G .

This theorem is proved by showing that both sides obey the same deletion-contraction relation and then checking the equation on the seed value of a graph with two vertices and a single edge between them whose matroid is $U_{1,1}$.

One can generalize the characteristic polynomial further by removing the restriction that it vanishes on loops. Tutte [96] introduced his eponymous polynomial for graphs by studying all possible topological invariants of graphs that were

well behaved under deletion and contraction. The definition was extended to matroids by Crapo [21]. We consider all invariants that are well behaved under deletion/contraction and multiplicative under direct sum. Let Mat be the class of all matroids. We define the Tutte–Grothendieck ring $K_0(\text{Mat})$ as the commutative ring given by formal sums of isomorphism classes of matroids subject to the relations

1. if i is neither a loop nor a coloop of a matroid M , then $[M] = [M \setminus i] + [M/i]$,
2. $[M_1 \oplus M_2] = [M_1][M_2]$,

This Tutte–Grothendieck ring was used implicitly by Tutte for graphs before Grothendieck’s introduction of the Grothendieck group. Because deletion and contraction commute, a matroid can be written uniquely in $K_0(\text{Mat})$ as a polynomial in the rank 1 matroids. Consequently, the Grothendieck group $K_0(\text{Mat})$ is the free polynomial ring over \mathbb{Z} generated by $U_{0,1}$ (a loop) and $U_{1,1}$ (a coloop).

Definition 7.2. Let R be a commutative ring. A Tutte–Grothendieck invariant valued in R is a homomorphism

$$h : K_0(\text{Mat}) \rightarrow R.$$

Definition 7.3. The rank generating polynomial of a matroid M is defined by

$$R_M(u, v) = \sum_{I \subseteq E} u^{r(M)-r(I)} v^{|I|-r(I)}.$$

The Tutte polynomial of M is defined by

$$T_M(x, y) = R_M(x - 1, y - 1).$$

The definition is justified by the following:

Proposition 7.4. The Tutte polynomial is the unique Tutte–Grothendieck invariant $T : K_0(\text{Mat}) \rightarrow \mathbb{Z}[x, y]$ satisfying $T_{U_{1,1}}(x, y) = x$ and $T_{U_{0,1}}(x, y) = y$.

This result naturally generalizes a formula for the characteristic polynomial that we will explore later.

This theorem shows that $U_{0,1}$ and $U_{1,1}$ are algebraically independent in $K_0(\text{Mat})$. The Tutte polynomial specializes to the characteristic polynomial:

$$\chi_M(q) = (-1)^{r(M)} T_M(1 - q, 0).$$

This can be seen by observing that the characteristic polynomial has the following properties: it takes the value $q - 1$ on coloops; it vanishes on loops; and it obeys a deletion/contraction relation. Its deletion/contraction relation has a sign which is accounted for by the fact that $r(M/i) = r(M) - 1$ for a non-coloop i .

The Tutte polynomial is well behaved under duality: $T_M(x, y) = T_{M^*}(y, x)$.

7.2 Motivic Definition of Characteristic Polynomial

Consider the Grothendieck ring of varieties over \mathbf{k} , $K_0(\text{Var}_{\mathbf{k}})$ where $\mathbf{k} = \mathbb{C}$ or $\mathbf{k} = \mathbb{F}_p$. This is the ring of formal sums of varieties where addition is disjoint union and multiplication is Cartesian product subject to the scissors relation: for closed $Z \subset X$,

$$[X] = [Z] + [X \setminus Z].$$

The identity of the ring is the class of a point.

One can show that there is a homomorphism

$$e : K_0(\text{Var}_{\mathbf{k}}) \rightarrow \mathbb{Z}[q]$$

such that

$$e([\mathbf{k}]) = q.$$

Definition 7.5. For $V^{d+1} \subseteq \mathbf{k}^{n+1}$, a $(d + 1)$ -dimensional linear subspace, the *characteristic polynomial* of V is

$$\chi_V(q) \equiv e([V \cap (\mathbf{k}^*)^{n+1}]).$$

Example 7.6. By inclusion/exclusion, for a generic subspace, we have

$$\chi_V(q) = q^{d+1} - \binom{d+1}{1}q^r + \binom{d+1}{2}q^{d-1} - \dots + (-1)^{d+1} \binom{d+1}{0}.$$

Here, we begin with V remove the $d + 1$ coordinate hyperplanes, add back in $\binom{d+1}{2}$ codimension 2 coordinate flats, and so on.

Theorem 7.7. Let $V \subseteq \mathbf{k}^{n+1}$ be a linear subspace with matroid M . Then, we have the following equality in $K_0(\text{Var}_{\mathbf{k}})$:

$$[V \cap (\mathbf{k}^*)^{n+1}] = \chi_M([k]).$$

Consequently, by applying e , we obtain the equality $\chi_V(q) = \chi_M(q)$.

Proof. We induct on n . For $n = 0$, if V is disjoint from \mathbf{k}^* , then it is contained in the point $x_0 = 0$. Consequently, 0 is a loop in M and both sides of the equality are 0. If V intersects \mathbf{k}^* , then $V = \mathbf{k}$ and $[V \cap \mathbf{k}^*] = [\mathbf{k}] - 1$. Then $M = U_{1,1}$, and $\chi_M([\mathbf{k}]) = [\mathbf{k}] - 1$.

In general, if M has a loop, then V is contained in a coordinate subspace, and each side of the equation is 0. Therefore, we may suppose that M is loopless.

If $M = M_1 \oplus M_2$, then $V = V_1 \oplus V_2 \subseteq \mathbf{k}^{n_1} \oplus \mathbf{k}^{n_2}$. Consequently,

$$[V \cap (\mathbf{k}^*)^{n+1}] = [V_1 \cap (\mathbf{k}^*)^{n_1}][V_2 \cap (\mathbf{k}^*)^{n_2}] = \chi_{M_1}([\mathbf{k}])\chi_{M_2}([\mathbf{k}])$$

by induction.

Therefore, we may suppose that M is a connected matroid and, therefore, has neither loops nor coloops. Pick $i \in E$, and let $H_i \subset \mathbf{k}^{n+1}$ be the hyperplane cut out by $x_i = 0$. Let A_i be the set given by

$$A_i = \{(x_0, x_1, \dots, x_n) \mid x_j \neq 0 \text{ if } i \neq j\}.$$

Then, we have the following motivic equation

$$[V \cap (\mathbf{k}^*)^{n+1}] = [V \cap A_i] - [V \cap A_i \cap H_i].$$

Since i is not a coloop, the coordinate projection $\pi : \mathbf{k}^{n+1} \rightarrow \mathbf{k}^n$ forgetting the i th component is injective on V . Moreover, π takes $V \cap A_i$ bijectively to $\pi(V) \cap (\mathbf{k}^*)^n$. Now, $\pi(V) \subseteq \mathbf{k}^n$ is the subspace representing $M \setminus i$. On the other hand, $V \cap A_i \cap H_i$ is the intersection of $V \cap H_i$ with the algebraic torus in H_i . Now, $V \cap H_i$ is the subspace representing M/i . Therefore, by induction, we have

$$[V \cap (\mathbf{k}^*)^{n+1}] = \chi_{M \setminus i}([\mathbf{k}]) - \chi_{M/i}([\mathbf{k}]) = \chi_M([\mathbf{k}]).$$

□

Note that $\chi_M(q)$ has degree $d + 1$ with the leading coefficient equal to 1 as the only $d + 1$ -dimensional variety that contributes to the motivic expression is V itself.

We now relate the characteristic polynomial to geometric invariants of $V \cap (\mathbf{k}^*)^{n+1}$. If $V \subset \mathbf{k}^{n+1}$ is a subspace, the motivic class of $[V \cap (\mathbf{k}^*)^{n+1}]$ lies in the commutative subring with unity $\langle [\mathbf{k}] \rangle \subseteq K_0(\text{Var}_{\mathbf{k}})$ generated by $[\mathbf{k}]$. Consequently, for any ring R and homomorphism $e' : K_0(\text{Var}_{\mathbf{k}}) \rightarrow R$, $e'([V \cap (\mathbf{k}^*)^{n+1}])$ is determined by $e'([\mathbf{k}]) \in R$. It follows that

$$e'([V \cap (\mathbf{k}^*)^{n+1}]) = \chi_V(e'[\mathbf{k}]).$$

If $\mathbf{k} = \mathbb{C}$, we can choose the homomorphism $\chi_y^c : K_0(\text{Var}_{\mathbf{k}}) \rightarrow \mathbb{Z}[u]$ to be the compactly supported χ_y -characteristic given on varieties by

$$\chi_y^c(X) = \sum_{p,q} \sum_m (-1)^m h^{p,q}(H_c^m(X)) u^p$$

where the Hodge numbers are taken with respect to Deligne’s mixed Hodge structure [76]. Since $\chi_y^c([\mathbf{k}]) = u$, we derive the following result related to that of Orlik–Solomon [72] as proved in [2, 54]

Theorem 7.8. *Let $V \subseteq \mathbb{C}^{n+1}$ be a subspace, then we have*

$$\chi_y^c([V \cap (\mathbb{C}^*)^{n+1}]) = \chi_V(u).$$

Similarly, if $\mathbf{k} = \mathbb{F}_{p^r}$, counting the number of \mathbb{F}_{p^r} -points gives a homomorphism $K_0(\text{Var}_{\mathbf{k}}) \rightarrow \mathbb{Z}$. Since $|\mathbb{F}_{p^r}| = p^r$, we have the following theorem of Athanasiadis [6]:

Theorem 7.9. *Let $V \subseteq \mathbb{F}_{p^r}^{n+1}$ be a subspace, then we have*

$$|V \cap (\mathbb{F}_{p^r}^*)^{n+1}| = \chi_V(p^r).$$

7.3 Möbius Inversion and the Characteristic Polynomial

The characteristic polynomial has a description given by inclusion/exclusion along the poset of flats. This is phrased in the language of Möbius inversion which is a combinatorial abstraction of the above motivic arguments. We will only make use of the Möbius function evaluated from the minimal flat. Let $L(M)$ be the poset of flats of M ordered by inclusion. We define $\mu(E, F)$ for every pair of flats $E \subseteq F$ recursively by

1. $\mu(E, E) = 1$,
2. $\mu(E, F) = -\sum_{E \subseteq F' \subsetneq F} \mu(E, F')$.

We will suppose that M is loopless and concentrate on $\mu(\emptyset, F)$. For $I \subseteq E = \{0, 1, \dots, n\}$, let L_I be the coordinate subspace

$$L_I = \{x \mid x_i = 0 \text{ if } i \in I\}.$$

Note that $\text{codim}(V \cap L_I) = r(I)$ and $V \cap L_I = V \cap L_F$ where F is the smallest flat containing I . Let

$$L_I^* = \{x \mid x_i = 0 \text{ if and only if } i \in I\}.$$

The following lemma, which is a special case of Möbius inversion, is the motivation for the above definition.

Lemma 7.10. *We have the following equality in $K_0(\text{Var}_{\mathbf{k}})$:*

$$[V \cap (\mathbf{k}^*)^{n+1}] = \sum_{F \in L(M)} \mu(\emptyset, F)[V \cap L_F].$$

Proof. We note that for any flat F ,

$$[V \cap L_F] = \sum_{\substack{F' \in L(M) \\ \text{mathbb{F}F' \supseteq F}} [V \cap L_{F'}^*].$$

Consequently,

$$\begin{aligned} \sum_{F \in L(M)} \mu(\emptyset, F)[V \cap L_F] &= \sum_{\substack{F, F' \in L(M) \\ \mathbb{1}F \subseteq F'}} \mu(\emptyset, F)[V \cap L_{F'}^*] \\ &= \sum_{F' \in L(M)} \left(\sum_{\substack{F \in L(M) \\ \mathbb{1}F \subseteq F'}} \mu(\emptyset, F) \right) [V \cap L_{F'}^*] \\ &= [V \cap L_{\emptyset}^*]. \end{aligned}$$

□

As a consequence of Weisner’s theorem [91, Sect. 3.9], [92, Theorem 3.10] for loopless matroids, we have the following result which will be important in the sequel:

Lemma 7.11. *Let F be a flat of a loopless matroid M . For any $a \in F$,*

$$\mu(\emptyset, F) = - \sum_{a \notin F' \triangleleft F} \mu(\emptyset, F')$$

where $F' \triangleleft F$ means that $F' \subset F$ and $r(F') = r(F) - 1$.

We have the following description of the characteristic polynomial:

Theorem 7.12. *The characteristic polynomial for a loopless matroid M of rank $d + 1$ is given by*

$$\chi_M(q) = \sum_{F \in L(M)} \mu(\emptyset, F)q^{d+1-r(F)}.$$

If M has loops, then $\chi_M(q) = 0$.

This theorem can be proven for representable matroids by combining Lemma 7.10 and Theorem 7.7. Before we give the proof of this theorem, we will need the following lemma:

Lemma 7.13. *Let F be a flat of a matroid M . Then*

$$U_F := \sum_{\substack{S \subseteq E \\ \text{cl}(S) = F}} (-1)^{|S|} = \begin{cases} \mu(\emptyset, F) & \text{if } M \text{ is loopless} \\ 0 & \text{otherwise.} \end{cases}$$

Proof. First, suppose M is loopless. We will show that U_F agrees with $\mu(\emptyset, F)$ inductively. It is clear that $U_\emptyset = 1$ and so it suffices to show (where \leq means containment of flats)

$$U_F = - \sum_{\emptyset \leq F' < F} U_{F'}.$$

Now,

$$\sum_{\emptyset \leq F' \leq F} U_{F'} = \sum_{\emptyset \leq F' \leq F} \sum_{\substack{S \subseteq E \\ \text{cl}(S) = F'}} (-1)^{|S|} = \sum_{S \subseteq F} (-1)^{|S|} = (1 - 1)^{|S|} = 0.$$

Suppose that M has loops. Let $F \subseteq E$ be the minimal flat of M , which consists of all the loops. Then

$$U_F = \sum_{S \subseteq F} (-1)^{|S|} = (1 - 1)^{|S|} = 0.$$

The same inductive argument shows that $U_F = 0$ for all flats. □

Note that in the notation of the above proof, Theorem 7.12 is equivalent to

$$\chi_M(q) = \sum_{F \in L(M)} U_F q^{d+1-r(F)}.$$

Now we continue with the proof of Theorem 7.12.

Proof. We verify the axioms for the characteristic polynomial.

By definition, the formula is true if M has a loop.

If $M = U_{1,1}$, it has two flats. We have $\mu(\emptyset, \emptyset) = 1, \mu(\emptyset, \{0\}) = -1$. Consequently, our formula gives $q - 1$.

Suppose $M = M_1 \oplus M_2$. If either M_i has a loop, then so does M and the formula is verified. Otherwise, the lattice of flats obeys $L(M) = L(M_1) \times L(M_2)$. It is easily seen that for $F_1 \in L(M_1), F_2 \in L(M_2)$, the Möbius function obeys

$$\mu_M(\emptyset, F_1 \times F_2) = \mu_{M_1}(\emptyset, F_1) \mu_{M_2}(\emptyset, F_2).$$

Therefore, $\chi_M(q) = \chi_{M_1}(q) \chi_{M_2}(q)$.

Now, we verify the deletion/contraction relation using an approach due to Whitney. Let $i \in E$ be neither a loop nor coloop. Then we have

$$\begin{aligned}
 \sum_{F \in L(M)} U_F q^{d+1-r(F)} &= \sum_{F \in L(M)} \left(\sum_{\substack{S \subseteq F \\ \text{cl}(S) = F}} (-1)^{|S|} \right) q^{d+1-r(F)} \\
 &= \sum_{S \subseteq E} (-1)^{|S|} q^{d+1-r(S)} \\
 &= \sum_{S \subseteq E \setminus \{i\}} (-1)^{|S|} q^{d+1-r(S)} + \sum_{S \subseteq E \setminus \{i\}} (-1)^{|S \cup \{i\}|} q^{d+1-r(S \cup \{i\})} \\
 &= \sum_{S \subseteq E \setminus \{i\}} (-1)^{|S|} q^{d+1-r_{M \setminus i}(S)} - \sum_{S \subseteq E \setminus \{i\}} (-1)^{|S|} q^{d-r_{M/i}(S)} \\
 &= \sum_{F \in L(M \setminus i)} (U_{M \setminus i})_F q^{d+1-r_{M \setminus i}(F)} - \sum_{F \in L(M/i)} (U_{M/i})_F q^{d-r_{M/i}(S)} \\
 &= \chi_{M \setminus i}(q) - \chi_{M/i}(q)
 \end{aligned}$$

where we have used the fact that because i is neither a loop nor a coloop,

$$r(M \setminus i) = d + 1, \quad r(M/i) = d.$$

□

We see from Lemma 7.11 that the coefficients of the characteristic polynomial alternate and the degree of the characteristic polynomial is equal to the rank of the matroid.

7.4 The Reduced Characteristic Polynomial

In this section, we introduce the reduced characteristic polynomial which will be important for the proof of the log-concavity of the characteristic polynomial of representable matroids. We will only consider loopless matroids. We begin by noting that $\chi_M(q)$ is divisible by $q - 1$. This can be proved combinatorially. In the representable case, we interpret $\chi_M(q)$ as the image of the motivic class of $V \cap (\mathbf{k}^*)^{n+1}$. Because \mathbf{k}^* acts freely on $V \cap (\mathbf{k}^*)^{n+1}$ we have

$$e([V \cap (\mathbf{k}^*)^{n+1}]) = e((([V \cap (\mathbf{k}^*)^{n+1}]/\mathbf{k}^*))e([\mathbf{k}^*])$$

with $e([\mathbf{k}^*]) = q - 1$.

Definition 7.14. The reduced characteristic polynomial is

$$\bar{\chi}_M(q) = \frac{\chi_M(q)}{q-1}.$$

Define the numbers $\mu^0, \mu^1, \dots, \mu^d$ by

$$\bar{\chi}_M(q) = \sum_{i=0}^d (-1)^i \mu^i q^{d-i}.$$

Note that $\mu^0 = 1$. We follow the convention that $\mu^{-1} = 0$.

Lemma 7.15. Let $a \in E$. The coefficients of $\bar{\chi}_M(q)$ are given by

$$\mu^k = (-1)^k \sum_{\substack{F \in L(M)_k \\ a \notin F}} \mu(\emptyset, F) = (-1)^{k+1} \sum_{\substack{F \in L(M)_{k+1} \\ a \in F}} \mu(\emptyset, F)$$

where the sum over the rank k flats not containing a .

Proof. We begin by proving the second equality by applying Lemma 7.11

$$\begin{aligned} \sum_{\substack{F \in L(M)_{k+1} \\ a \in F}} \mu(\emptyset, F) &= - \sum_{\substack{F \in L(M)_{k+1} \\ a \in F}} \sum_{\substack{F' \subsetneq F \\ a \notin F'}} \mu(\emptyset, F') \\ &= - \sum_{\substack{F' \in L(M)_k \\ a \notin F'}} \sum_{\substack{F \supsetneq F' \\ a \in F}} \mu(\emptyset, F') \\ &= - \sum_{\substack{F' \in L(M)_k \\ a \notin F'}} \mu(\emptyset, F') \end{aligned}$$

where the last equality follows from the fact that for F' not containing a , there is a unique flat F with $r(F) = r(F') + 1$ and $a \in F$. If we write

$$\chi_M(q) = \mu_0 q^{r+1} - \mu_1 q^r + \dots + (-1)^{d+1} \mu_{r+1}$$

then

$$\mu_k = \sum_{F \in L(M)_k} \mu(\emptyset, F)$$

and by equating coefficients of $\bar{\chi}_M(q)$ and $\frac{\chi_M(q)}{q-1}$,

$$\mu_k = \mu^k + \mu^{k-1}.$$

The theorem is true for $k = 0$. It now follows by induction using the above formula. \square

7.5 Log-Concavity of Whitney Numbers

We now discuss some conjectured inequalities for characteristic polynomials.

Definition 7.16. A polynomial

$$\chi(q) = \mu_0 q^{r+1} + \mu_1 q^r + \cdots + \mu_{r+1}$$

is said to be *unimodal* if the coefficients are unimodal in absolute value, i.e. there is a j such that

$$|\mu_0| \leq |\mu_1| \leq \cdots \leq |\mu_j| \geq |\mu_{j+1}| \geq \cdots \geq |\mu_{r+1}|.$$

The chromatic polynomial of a graph was conjectured to be unimodal by Read [79]. This conjecture was proved by Huh in 2010 [44]. In fact, Huh did more. He proved the characteristic polynomial of a matroid representable over a field of characteristic 0 is log-concave.

Definition 7.17. The polynomial $\chi(q)$ is said to be *log-concave* if for all i ,

$$|\mu_{i-1}\mu_{i+1}| \leq \mu_i^2.$$

If the inequality is strict, then the polynomial is said to be strictly log-concave. Log-concavity means that the logarithms of the absolute values of the (alternating) coefficients form a concave sequence. Log-concavity implies unimodality as long as the sequence of coefficients has no internal zeroes. Below we will give a streamlined version of the proof of Huh and the author [47] that the characteristic polynomial of any representable matroid is log-concave. This proof is very similar to Huh’s original proof except that it replaces singularity theory with intersection theory on toric varieties. This theorem is a special case of the still-open Rota–Heron–Welsh conjecture:

Conjecture 7.18. For any matroid, $\chi_M(q)$ is log-concave.

The characteristic polynomial of a matroid representable over a field of characteristic 0 was proved to be strictly log-concave by Huh in [46] by studying the variety of critical points that arises in maximum likelihood estimation.

A closely related conjecture is Mason’s conjecture. Let M be a matroid of rank $d + 1$. Let \mathcal{I} be the set of independent sets of M . We define the f -vector by setting f_i to be the number of independent sets of cardinality i . It was conjectured by Mason the f -vector was log-concave, that is the polynomial

$$f(q) = \sum_{i=0} f_i q^{r+1-i}$$

is log-concave. In fact, Mason strengthened his conjecture to these statements for every k to the following:

1. $f_k^2 \geq \frac{k+1}{k} f_{k-1} f_{k+1}$
2. $f_k^2 \geq \frac{k+1}{k} \frac{f_{1-k+1}}{f_{1-k}} f_{k-1} f_{k+1}$

Mason’s original log-concavity conjecture was proved in the representable case by Lenz:

Theorem 7.19 ([64]). *The f -vector of a realizable matroid is strictly log concave.*

The h -polynomial is defined by

$$h(q) = f(q - 1).$$

Dawson [24] conjectures this to be log-concave. Huh [46] proved that h is log-concave and without internal zeroes if M is representable over a field of characteristic 0.

A much less tractable sequence of numbers conjectured to be log-concave are the Whitney numbers of the second kind. They are W_k , the number of flats of M of rank k . There are strengthenings of the log-concavity conjecture for W_k by Mason analogous to those of the f -vector. A special case is the Points–Lines–Planes conjecture:

Conjecture 7.20. If M is a rank four matroid, then

$$W_2^2 \geq \frac{3}{2} W_1 W_3.$$

In the representable case, we imagine having W_1 points in 3-space determining W_2 lines and W_3 planes. A slightly strengthening of this conjecture was proved by Seymour for all matroids in which no line has more than four points [86]. See [1] for a survey.

8 Matroid Polytopes

8.1 Definition of Matroid Polytope

We first give a historical introduction to matroid polytopes using the independence polytope of Edmonds [32], which he introduced as a tool in combinatorial optimization. Specifically, Edmonds was interested in maximizing a weight function along independent subsets. Let $c : E \rightarrow \mathbb{R}$ be a function on the ground set. For $I \subseteq E$, an independent subset, set the weight function on I to be

$$c(I) = \sum_{i \in I} c(i).$$

Now, a natural problem is to find a basis for the matroid of maximum weight. Matroids have the property that this problem is solvable by a greedy algorithm:

1. Set $B := \emptyset$,
2. While B is not a basis, pick $i \notin B$ such that $B \cup \{i\}$ is independent and i is of maximum weight and replace B by $B \cup \{i\}$.

The resulting set B is a basis of maximum weight. Moreover, matroids can be characterized as a structure for which the greedy algorithm succeeds.

The basis optimization problem can also be phrased in terms of maximizing a linear functional on the independence polytope $P(\mathcal{I}(M)) \subset \mathbb{R}^E$ given by

$$P(\mathcal{I}(M)) = \text{Conv}(\{e_I \mid I \in \mathcal{I}\})$$

where e_I is the characteristic vector of an independent set I ,

$$e_I = \sum_{i \in I} e_i.$$

We can define a *weight vector* $w \in \mathbb{R}^n$ by

$$w = \sum c(i)e_i.$$

The set of points $x \in P$ that maximize the function $\langle x, w \rangle$ is a face of the independence polytope. The vertices of that faces are the independent sets of maximum weight.

Edmonds studied the structure of independence polytope and was able to prove the following important theorem in combinatorial optimization:

Theorem 8.1. *Let M_1, M_2 be two matroids on the same ground set where intersection is given by intersecting the collection of independent subsets. Then the independence polytope is well behaved under intersections in the following sense:*

$$P(\mathcal{I}(M_1) \cap \mathcal{I}(M_2)) = P(\mathcal{I}(M_1)) \cap P(\mathcal{I}(M_2))$$

This theorem has relevance to finding common transversals. We recommend [13] as a survey on matroid polytopes in optimization.

Matroid polytopes received renewed attention due to work of Gelfand–Goresky–MacPherson–Serganova [42]. In this case, they investigated the convex hull of the characteristic vectors of the bases of a matroid.

Definition 8.2. The *matroid polytope* of a matroid M is given as the following convex hull:

$$P(M) = \text{Conv}(\{e_B \mid B \in \mathcal{B}\}).$$

Lemma 8.3. *The vertices of the matroid polytope $P(M)$ are exactly the characteristic vectors of the bases of M .*

Proof. If e_B was a convex combination of $\{e_{B_1}, \dots, e_{B_k}\}$, then we could write

$$e_B = \sum_j \lambda_j e_{B_j}$$

for $0 < \lambda_j, \sum \lambda_j = 1$. For each $i \in B$, by extracting the component of e_i from the above equation, we would have $i \in B_j$ for all j . Therefore, we conclude $B_1 = \dots = B_k = B$. Consequently, e_B cannot be written as a non-trivial convex combination. \square

Example 8.4. The matroid polytope of the uniform matroid $U_{d+1,n+1}$ is $\Delta(d+1, n+1)$, the $(d+1, n+1)$ -hypersimplex which is defined to be the convex hull of the $\binom{n+1}{d+1}$ vectors e_B as B ranges over all $(d+1)$ -element subsets of $E = \{0, 1, \dots, n\}$.

Example 8.5. The matroid polytope of $U_{2,4}$ is an octahedron in \mathbb{R}^4 whose vertices are

$$\{(1, 1, 0, 0), (1, 0, 0, 1), (1, 0, 1, 0), (0, 1, 1, 0), (0, 1, 0, 1), (0, 0, 1, 1)\}$$

where there is an edge between a pair of vertices if and only if they differ in exactly two coordinates.

Remarkably, matroid polytopes can be abstractly characterized leading to a cryptomorphic definition of matroids by the work of Gelfand–Goresky–MacPherson–Serganova:

Definition 8.6. A *matroid polytope* P is a convex lattice polytope in \mathbb{R}^{n+1} which is contained in $\Delta(d+1, n+1)$ for some d such that all vertices of P are vertices of $\Delta(d+1, n+1)$ and each edge is a translate of $e_i - e_j$ for $i \neq j$.

Proposition 8.7. *Let M be a rank $d+1$ matroid on $E = \{0, 1, \dots, n\}$. Then $P(M)$ is a matroid polytope.*

Proof. We sketch the proof. Clearly every vertex of $P(M)$ is a vertex of $\Delta(d+1, n+1)$. Let L be the line segment between e_{B_1} and e_{B_2} . If we have $|B_1 \Delta B_2| \neq 2$, we must show that L is not an edge of $P(M)$. By the repeated use of the basis exchange axiom, one can show that the midpoint of the edge $\frac{e_{B_1} + e_{B_2}}{2}$ lies in the convex hull of other vertices of the matroid polytope. \square

We give the proof of the converse following [42].

Theorem 8.8. *Let P be a matroid polytope. Then there exists a unique matroid M such that $P = P(M)$.*

Proof. Let \mathcal{B} be the set of $(d + 1)$ -element subsets of $E = \{0, 1, \dots, n\}$ such that $\{e_B \mid B \in \mathcal{B}\}$ is the set of vertices of P . We will show that \mathcal{B} is the collection of bases for a matroid by verifying the basis exchange axiom.

Let $B_1, B_2 \in \mathcal{B}$. Let $i \in B_1 \setminus B_2$. We must find $j \in B_2 \setminus B_1$ such that $(B_1 \setminus i) \cup j \in \mathcal{B}$. Let B'_1, \dots, B'_k be the bases such that $\{e_{B_1} e_{B'_l}\}_l$ are the edges of P at e_{B_1} . If B_2 is among the B'_l 's, then we are done because then $|B_1 \Delta B_2| = 2$. Otherwise, because the midpoint between e_{B_1} and e_{B_2} is contained in P , it is contained in the convex cone at e_{B_1} spanned by the edges containing e_{B_1} . Therefore, we may write

$$\frac{e_{B_1} + e_{B_2}}{2} = e_{B_1} + \sum_l \lambda_l (e_{B'_l} - e_{B_1})$$

or

$$\frac{e_{B_2} - e_{B_1}}{2} = \sum_l \lambda_l (e_{B'_l} - e_{B_1}) \tag{1}$$

for $\lambda_l \geq 0$. Extracting the coefficient of e_i , we get

$$-\frac{1}{2} = - \sum_{l \mid i \notin B'_l} \lambda_l.$$

Let l' be such that $i \notin B_{l'}$ and $\lambda_{l'} > 0$. Let j be the unique element of $B_{l'} \setminus B_1$. Because the coefficient of e_j on the right side of (1) is positive, we have that $j \in B_2 \setminus B_1$. \square

It turns out that the dimension of $P(M)$ is $n + 1 - \kappa(M)$ where $\kappa(M)$ is the number of connected components of M . In fact, if M is connected, then the dimension of $P(M)$ is n for the following reasons: all vertices of the matroid polytope lie in the hyperplane described by $x_0 + x_1 + \dots + x_n = d + 1$ corresponding to the fact that all of the bases have the same number of elements; and for all pairs $i \neq j$, $e_i - e_j$ is parallel to an edge of $P(M)$. If E is partitioned as $E = E_1 \sqcup \dots \sqcup E_k$ where $M = M_1 \oplus \dots \oplus M_k$ and M_i is a matroid on E_i , then $P(M) = P(M_1) \oplus \dots \oplus P(M_k) \subset \mathbb{R}^{E_1} \oplus \dots \oplus \mathbb{R}^{E_k}$.

8.2 Faces of the Matroid Polytope

In this section, we study faces of the matroid polytope following [3]. They correspond to bases maximizing a certain weight vector. Further details can be found in [33]. We first review faces of polytopes.

Definition 8.9. Let P be a convex polytope in \mathbb{R}^{n+1} . For $w \in \mathbb{R}^{n+1}$, let P_w be the set of points $x \in P$ that minimize the function $\langle x, w \rangle$.

The set P_w is a face of the polytope P . Observe that any face of the matroid polytope is itself a matroid polytope. We will identify its matroid.

Definition 8.10. Let M be a matroid on $E = \{0, \dots, n\}$ and let $w \in \mathbb{R}^E$. For a basis $B \in \mathcal{B}(M)$, let the w -weight of B be

$$w_B = \sum_{i \in B} w_i.$$

Let M_w be the matroid whose bases are the bases of M of minimal w -weight.

We know that M_w is a matroid for the following reason: the convex hull of the characteristic vectors of its bases is the matroid polytope $P(M)_w$, allowing us to apply Theorem 8.8.

We give a description of the matroid M_w following [3]. Let S_\bullet be the flag of subsets

$$S_\bullet = \{\emptyset = S_0 \subsetneq S_1 \subsetneq \dots \subsetneq S_k \subseteq F_{k+1} = E\},$$

where w is constant on $S_i \setminus S_{i-1}$ and $w|_{S_i \setminus S_{i-1}}$ is strictly increasing with i .

Lemma 8.11. *We have the following equality:*

$$M_w = \bigoplus_{i=1}^{k+1} (M|_{S_i})/S_{i-1}.$$

Proof. Every basis of minimal w weight can be constructed by the greedy algorithm. In others words, it will start with the empty set, add elements of S_1 as long as the result is independent, move on to S_2 , and so on. Consequently, a basis of M_w consists of $r(S_i) - r(S_{i-1})$ elements of $S_i \setminus S_{i-1}$ for all i . Such a set of $r(S_i) - r(S_{i-1})$ elements is exactly a basis of $(M|_{S_i})/S_{i-1}$. \square

We will make use of this description in Sect. 11 when we discuss the Bergman fan of a matroid.

8.3 Valuative Invariants

The polytope definition of matroids allows one to think about invariants of matroids in a different light. One can study invariants of matroids that are well behaved under subdivisions of the matroid polytope. This is the approach taken by Fink and Derksen [26, 27].

Definition 8.12. A *matroidal subdivision* of a matroid polytope $P(M) \subset \mathbb{R}^{n+1}$ is a polyhedral complex \mathcal{D} whose cells are polytopes in \mathbb{R}^{n+1} such that

- (1) Each cell of \mathcal{D} is a matroid polytope, and
- (2) $\bigcup_{P \in \mathcal{D}} P = P(M)$

Example 8.13. As an example of a matroid subdivision, consider $U_{2,4}$ whose matroid polytope is an octahedron. It has a matroid subdivision into two pyramids that meet along the square whose vertices are $\{(1, 0, 0, 1), (1, 0, 1, 0), (0, 1, 1, 0)\}, \{(0, 1, 0, 1)\}$ [89]. In fact, let M_{top} be the rank 2 matroid on $E = \{0, 1, 2, 3\}$ where 0 and 1 are taken to be parallel but all other two element subsets are bases. Its matroid polytope is the top pyramid whose vertices are $\{(0, 0, 1, 1), (1, 0, 0, 1), (1, 0, 1, 0), (0, 1, 1, 0), (0, 1, 0, 1)\}$. Similarly, the bottom pyramid is the matroid polytope corresponding to M_{bot} where 2 and 3 are mandated to be parallel. The middle square is the matroid polytope of M_{mid} where 0 and 1 are mandated to be parallel as are 2 and 3 while the bases are $\{0, 2\}, \{1, 2\}, \{0, 3\}, \{1, 3\}$.

Write P_1, \dots, P_k for the top-dimensional cells of \mathcal{D} . For $J \subseteq \{1, \dots, k\}$, let $P_J = \bigcap_{j \in J} P_j$. Let \mathcal{M} be the set of all matroid polytopes in \mathbb{R}^{n+1} and let A be an abelian group.

Definition 8.14. A function $f : \mathcal{M} \rightarrow A$ is a *matroid valuation* if for all matroidal subdivisions of a matroid polytope $P(M)$,

$$f(P(M)) = \sum_{\emptyset \neq J \subseteq \{1, \dots, k\}} (-1)^{|J|-1} f(P_J).$$

In other words, f obeys an inclusion/exclusion relation for matroidal subdivisions. Consequently, a valuative invariant of lattice polytopes like the Ehrhart polynomial immediately gives valuative invariants of matroids. When the subdivision is regular, the matroidal subdivision can be interpreted as a tropical linear subspace [89] which is a combinatorial abstraction of a family of linear subspaces defined over a disk that degenerates into a union of linear subspaces.

The Tutte polynomial along with a quasi-symmetric function-valued invariant introduced by Billera–Jia–Reiner [10] is valuative. Derksen [26] introduced a valuative invariant which was proved to be universal among valuative invariants [27]. It takes values in the ring of quasi-symmetric functions, whose underlying vector space has a basis indexed by compositions $\underline{\alpha} = (\alpha_1, \alpha_2, \dots, \alpha_l)$. Let $U_{\underline{\alpha}}$ be some basis for the ring of quasi-symmetric functions over \mathbb{Z} . Derksen’s invariant \mathcal{G} is defined as

$$\mathcal{G} = \sum_{\underline{E}} U_{r(\underline{E})}$$

where $\underline{E} = \{\emptyset = E_0 \subset E_1 \subset E_2 \subset \dots \subset E_{n+1} = E\}$ ranges over all $(n + 1)!$ maximal chains of subsets of E and

$$r(\underline{E}) = (r(E_1) - r(E_0) + 1, r(E_2) - r(E_1) + 1, \dots, r(E_{n+1}) - r(E_n) + 1).$$

This invariant specializes to any valutive invariant under homomorphisms from quasi-symmetric functions to an abelian group [27]. This is proved by showing that Schubert matroids (after allowing permutations of the ground set) are a basis for the dual space to valutive invariants.

9 Grassmannians and Matroid Polytopes

9.1 Plücker Coordinates and Basis Exchange

The Grassmannian $\text{Gr}(d + 1, n + 1)$ is an algebraic variety whose points correspond to the $(d + 1)$ -dimensional linear subspaces of \mathbf{k}^{n+1} . Equivalently, it is the set of all d -dimensional projective subspaces of \mathbb{P}^n . It generalizes projective space in the sense that $\text{Gr}(1, n + 1)$ is isomorphic to \mathbb{P}^n .

We present the Grassmannian as a quotient of the Stiefel variety following [43]. The Stiefel variety $\text{St}(d + 1, n + 1)$ parameterizes linearly independent $(d + 1)$ -tuples of vector in \mathbf{k}^{n+1} . We denote their span as the $(d + 1)$ -dimensional subspace $P \subset \mathbf{k}^{n+1}$ and write the $(d + 1)$ -tuple (v_1, \dots, v_n) as a $(d + 1) \times (n + 1)$ -matrix X . Now, Gl_{d+1} acts on the $(d + 1)$ -tuple without changing P . This is the same thing as the left-action of Gl_{d+1} on the matrix. Because the action is free and acts transitively on bases of P , the Grassmannian is the quotient

$$\text{Gr}(d + 1, n + 1) = \text{St}(d + 1, n + 1) / \text{Gl}_{d+1} .$$

There are homogeneous coordinates, the Plücker coordinates that embed this quotient into projective space: let B be a $(d + 1)$ -tuple of distinct elements of $\{1, \dots, n + 1\}$, the Plücker coordinate p_B is the determinant of the $(d + 1) \times (d + 1)$ -matrix formed by the columns indexed by B . The Plücker coordinates are indeed coordinates on the Grassmannian. Indeed, because X is rank $d + 1$, we may permute the columns of M so that the first $d + 1$ columns form a non-singular matrix. By applying an element of $\text{Gl}(d + 1)$, we may suppose X is the form

$$X = [I_{d+1} | A]$$

where A is a $(d + 1) \times (n - d)$ -matrix. Let B consist of d elements of $\{1, \dots, d + 1\}$ together with an element i of $\{d + 2, \dots, n + 1\}$. Then p_B is equal (up to sign) to an entry of the i th column of X . Therefore, from the Plücker coordinates, we can recover X .

Now, because if B and B' are related by a permutation, p_B and $p_{B'}$ are equal up to a sign, it suffices to consider only the Plücker coordinates corresponding to $B = (i_1, \dots, i_{d+1})$ with $i_1 < i_2 < \dots < i_{d+1}$. Therefore, there are $\binom{n+1}{d+1}$ Plücker coordinates. The Plücker coordinates give an embedding

$$\text{Gr}(d + 1, n + 1) \hookrightarrow \mathbb{P}^{\binom{n+1}{d+1}-1} .$$

The Plücker coordinates obey natural quadratic relations, called the Plücker relations: for $1 \leq i_1 < \dots < i_d \leq n + 1, 1 \leq j_1 < \dots < j_{d+2} \leq n + 1$, we have, for every point of the Grassmannian,

$$\sum_{a=1}^{d+1} (-1)^a p_{i_1 \dots i_d j_a} p_{j_1 \dots \hat{j}_a \dots j_{d+2}} = 0$$

where \hat{j}_a means j_a is deleted. Moreover, these relations generate the ideal defining $\text{Gr}(d + 1, n + 1)$.

Given a point of the Grassmannian whose Plücker coordinates are (p_B) , we can define a matroid by setting the bases to be

$$\mathcal{B} = \left\{ B \in \binom{\{1, \dots, n + 1\}}{d + 1} \mid p_B \neq 0 \right\}.$$

Lemma 9.1. *The set \mathcal{B} defined above is the set of bases for a matroid.*

Proof. Because the matrix X is of rank $d + 1$, the set \mathcal{B} is nonempty. We need only show that it satisfies the basis exchange axiom. Let $B_1, B_2 \in \mathcal{B}, i \in B_1 \setminus B_2$. We must show that there is an element $j \in B_2 \setminus B_1$ such that $(B_1 \setminus i) \cup \{j\} \in \mathcal{B}$. The left side of the Plücker relation for this choice of i 's and j 's has $p_{B_1} p_{B_2}$ as a summand. Because this monomial is non-zero, there must be another non-zero monomial of the form $p_{(B_1 \setminus \{i\}) \cup \{j\}} p_{(B_2 \setminus \{j\}) \cup \{i\}}$. □

We can view this matroid as coming from a vector configuration. Specifically, pick a matrix $X \in \text{St}(d + 1, n + 1)$ representing that linear subspace. Take the vectors to be the columns of X . The set of $d + 1$ vectors is a basis for the matroid if and only if the corresponding determinant is non-zero. This occurs exactly when these vectors give a basis of \mathbf{k}^{d+1} .

We note that $\text{Gr}(d + 1, n + 1)$ is isomorphic to $\text{Gr}(n - d, n + 1)$. Specifically if $V \subset \mathbf{k}^{n+1}$ is a d -dimensional subspace, there is a map dual to inclusion $i^* : (\mathbf{k}^{n+1})^* \rightarrow V^*$. The map is surjective so its kernel is $(n - d)$ -dimensional. This isomorphism is analogous to matroid duality.

9.2 Projective Toric Varieties

We give a review of not-necessarily-normal (or Sturmfeldian) projective toric varieties. They are important for relating the matroid polytope to the Grassmannian. For more details, see [19]. Let $\mathcal{A} \subset \mathbb{Z}^{n+1}$ be a finite subset. Write $\mathcal{A} = \{\omega_0, \omega_1, \dots, \omega_N\}$. We can define a morphism $i : (\mathbf{k}^*)^{n+1} \rightarrow \mathbb{P}^N$ by

$$x \mapsto [x^{\omega_0} : \dots : x^{\omega_N}]$$

where if $x = (x_0, \dots, x_n)$ and $\omega_i = (\omega_{i0}, \dots, \omega_{in})$,

$$x^{\omega_i} = x_0^{\omega_{i0}} \dots x_n^{\omega_{in}}.$$

Let A be the $(n + 1) \times (N + 1)$ -matrix whose columns are the ω 's. If A is unimodular and full rank, then i is an inclusion. The toric variety $\mathbb{P}_{\mathcal{A}}$ associated with \mathcal{A} is the closure $\overline{i((\mathbf{k}^*)^{n+1})}$. The weight polytope of $X_{\mathcal{A}}$ is the convex hull of \mathcal{A} . The torus orbits in $\mathbb{P}_{\mathcal{A}}$ are in bijective correspondence with the faces of the weight polytope.

This definition can be extended to include closures of torus orbits in projective spaces. Suppose that \mathbf{k} is an algebraically closed field. Let $H = (\mathbf{k}^*)^{n+1}$ act on a vector space V . There is a character decomposition

$$V = \bigoplus_{\chi} V_{\chi}$$

where for a character χ of H ,

$$V_{\chi} = \{v \in V \mid t \cdot v = \chi(t)v \text{ for all } t \in H\}.$$

Let $v \in \mathbb{P}(V)$. We will study the H -orbit closure $\overline{H \cdot v} \subseteq \mathbb{P}(V)$. Lift v to an element of $\tilde{v} \in V$, and write

$$\tilde{v} = \sum_{\chi} \tilde{v}_{\chi}$$

for $\tilde{v}_{\chi} \in V_{\chi}$. Let $\mathcal{A}(v) = \{\chi \mid \tilde{v}_{\chi} \neq 0\}$. By unpacking definitions, we have $\overline{H \cdot v} \cong X_{\mathcal{A}(v)}$. The *weight polytope* of v is the convex hull of $\mathcal{A}(v)$ in $N^{\vee} \otimes \mathbb{R}$ where N^{\vee} is the character lattice of H . If the image of v in $\mathbb{P}(V)$ is stabilized by a subtorus $H' \subseteq H$, there is an induced linear map $N^{\vee} \rightarrow (N')^{\vee}$ and the weight polytope lies in a translate of $\ker(N^{\vee} \rightarrow (N')^{\vee})$.

9.3 Torus Orbits and Matroid Polytopes

In this section, we consider a torus acting on the Grassmannian. Note that the algebraic torus $(\mathbf{k}^*)^{n+1}$ acts on \mathbb{P}^n by dilating the homogeneous coordinates. This induces an action on $\text{Gr}(d + 1, n + 1)$: for $t \in (\mathbf{k}^*)^{n+1}$, $P \subset \mathbb{P}^n$, a d -dimensional subspace, $t \cdot P$ is also a d -dimensional subspace. We can see this action in terms of the Stiefel variety: $(\mathbf{k}^*)^{n+1}$ dilates the columns of the matrix X . Note that the diagonal \mathbf{k}^* acts trivially on the Grassmannian since it preserves any subspace P .

For $P \in \text{Gr}(n - d, n)$, we may consider the closure of the $(\mathbf{k}^*)^{n+1}$ -orbit of P ,

$$\mathbb{P}_P = \overline{(\mathbf{k}^*)^{n+1} \cdot P}.$$

As observed in [90, Proposition 12.1], this orbit closure is a normal toric variety by arguments of White [104]. Now, by Lemma 9.1, we may produce a rank $d + 1$ matroid M on $\{0, 1, \dots, n\}$ from P .

We have the following combinatorial description of the toric variety by Gelfand–Goresky–MacPherson–Serganova [42]:

Theorem 9.2. *The weight polytope of \mathbb{P}_P is the matroid polytope $P(M)$.*

Proof. The $(\mathbf{k}^*)^{n+1}$ -action extends to the ambient space of the Plücker embedding $\mathbb{P}^{\binom{n+1}{d+1}-1}$. In fact, by the matrix description of the group action, if $t \in (\mathbf{k}^*)^{n+1}$ is given by (t_0, \dots, t_n) , then

$$t \cdot p_B = \left(\prod_{i \in B} t_i \right) p_B$$

Consequently, p_B lies in the character space given by e_B , the characteristic vector of B . The weight polytope is the convex hull of e_B for $B \in \mathcal{B}$. □

Note that the weight polytope lies in an affine hyperplane corresponding to the fact that all bases have $d + 1$ elements. This reflects the fact that the diagonal torus of $(\mathbf{k}^*)^{n+1}$ -stabilizes $P \in \text{Gr}(d + 1, n + 1)$.

9.4 Thin Schubert Cells

We can use matroids to define thin Schubert cells, locally closed subsets of the Grassmannian that refine Schubert cells. A thin Schubert cell consists of all subspaces with a given matroid. Thin Schubert cells were introduced in [42] by Gelfand–Goresky–MacPherson–Serganova. Prior to their introduction in the algebraic geometric context, White [103] studied their coordinate rings. Following the notation of [42], we will think of matroids as vector configuration given by the projection of the coordinate vectors.

The closure of the usual Schubert cells, called Schubert varieties, can be described by the vanishing of some Plücker coordinates. The data of the matroid corresponds to imposing exactly which Plücker coordinates vanish. We will use this data to define a thin Schubert cell.

Definition 9.3. Let M be a rank $d + 1$ matroid on $E = \{0, \dots, n\}$. The *thin Schubert cell* Gr_M in $\text{Gr}(d + 1, n + 1)$ is the locally closed subset given by

$$\text{Gr}_M = \{P \in \text{Gr}(d + 1, n + 1) \mid p_B(P) \neq 0 \text{ if and only if } B \text{ is a basis of } M\}.$$

Another convention is to consider the thin Schubert cell associated with M in $\text{Gr}(n - d, n + 1) = \text{Gr}(d + 1, n + 1)$ as follows: for an $(n - d)$ -dimensional subspace $P \in \text{Gr}(n - d, n + 1)$, there is a projection $\pi_P : \mathbf{k}^{n+1} \rightarrow \mathbf{k}^{n+1}/P$. For $S \subset \{0, \dots, n\}$,

let \mathbf{k}^S be the linear space spanned by e_i for $i \in S$. We may consider the matroid given by the vector configuration $\pi_P(e_0) \dots \pi_P(e_n)$ which has the rank function $r(I) = \dim(\pi_P(\mathbf{k}^I))$.

We can describe the stabilizer of Gr_M under the torus action in terms of M . Let $M = M_1 \oplus \dots \oplus M_\kappa$ be a connected component decomposition of M where each M_i is a matroid on F_i for a partition $E = F_1 \sqcup \dots \sqcup F_\kappa$. Then we have a decomposition of P ,

$$\mathbf{k}^{n+1} = \mathbf{k}^{F_1} \oplus \dots \oplus \mathbf{k}^{F_\kappa}$$

given by

$$P = P_1 \oplus \dots \oplus P_\kappa$$

for $P_i \subseteq \mathbf{k}^{E_i}$. We have an inclusion $(\mathbf{k}^*)^{F_i} \hookrightarrow \mathbf{k}^{n+1}$. The image of the diagonal torus in $(\mathbf{k}^*)^{F_i}$ stabilizes P_i . One can show that the stabilizer of P is $(\mathbf{k}^*)^\kappa$ corresponding to the direct product of these diagonal tori.

The thin Schubert cells do not provide a stratification of the Grassmannian. In fact, the closure of a thin Schubert cell is not always a finite union of thin Schubert cells. However, thin Schubert cells do occur as the intersection of finitely many Schubert cells with respect to different flags.

We have the following proposition which describes how the thin Schubert cells behave under direct sums of matroids:

Proposition 9.4. *If the matroid M has a unique decomposition into connected components*

$$M = M|_{F_0} \oplus \dots \oplus M|_{F_c},$$

for flats, F_i , then there is a natural isomorphism

$$\text{Gr}_M \cong \text{Gr}_{M|_{F_0}} \times \dots \times \text{Gr}_{M|_{F_c}}$$

9.5 Realization Spaces

Given a simple matroid M of rank $d + 1$ on $\{0, 1, \dots, n\}$, we can define *realization spaces*, space whose points correspond to representations of M . These spaces exist as moduli spaces, but we will focus our attention on their \mathbf{k} -points and treat them as parameter spaces. The results of this subsection, however, can be stated in moduli-theoretic language.

Thin Schubert cells can be related to a universal family of representations. Here, we say that a d -dimensional subspace $V \subset \mathbb{P}^n$ represents M if the matroid associated

with V is M . Recall that the Grassmannian $\text{Gr}(d + 1, n + 1)$ has a universal family of subspaces over it: there is a subvariety $\mathcal{V} \subset \mathbb{P}^n_{\text{Gr}(d+1,n+1)}$ that fits in a commutative diagram

$$\begin{array}{ccc}
 \mathcal{V} & \hookrightarrow & \mathbb{P}^n_{\text{Gr}(d+1,n+1)} \\
 \downarrow & & \swarrow \\
 \text{Gr}(d + 1, n + 1) & &
 \end{array}$$

such that \mathcal{V} is flat over $\text{Gr}(d + 1, n + 1)$ and the fiber \mathcal{V}_V for $V \in \text{Gr}(d + 1, n + 1)$ is the subspace of \mathbb{P}^n corresponding to V .

Definition 9.5. A family of representations of M over a scheme X is a flat family $V \rightarrow X$ of d -dimensional subspaces in \mathbb{P}^n_X such that for every closed point x , the fiber V_x is a representation of M .

The base-change $\mathcal{V}_{\text{Gr}^M}$ over the thin Schubert cell Gr^M is a family of realizations. Its \mathbf{k} -points, $\text{Gr}^M(\mathbf{k})$ are exactly the representations of M over \mathbf{k} .

One can also consider the space of representations of M as a hyperplane arrangement. Recall that if $V \subset \mathbb{P}^n$ is a representation of M as a subspace, intersecting V with each coordinate hyperplane gives an arrangement of $n + 1$ hyperplanes on $\mathbb{P}^d \cong V$ representing M . We can alternatively consider an arrangement of $n + 1$ hyperplanes on \mathbb{P}^d by considering an $(n + 1)$ -tuple of non-zero linear forms on \mathbb{P}^d up to constant multiples. Let $\Gamma = \mathbb{A}^{d+1} \setminus \{0\}$. Therefore, an $n + 1$ -tuple of non-zero linear forms (s_0, s_1, \dots, s_n) is described by a point of Γ^{n+1} .

Definition 9.6. The *linear form realization space* for M is the subscheme Γ_M of Γ^{n+1} consisting of hyperplane arrangements realizing M .

Two tuples of sections cut out the same family of hyperplane arrangements if and only if they differ by an element of $(\mathbf{k}^*)^{n+1}$ scaling the sections (s_0, s_1, \dots, s_n) and two families of hyperplane arrangements are isomorphic, by definition, if they differ by an element of PGL_{d+1} . These actions commute, and the product $\text{PGL}_{d+1} \times (\mathbf{k}^*)^{n+1}$ acts freely on Γ_M .

The quotient

$$\mathbf{R}_M = \Gamma_M / (\text{PGL}_{d+1} \times (\mathbf{k}^*)^{n+1})$$

is called the *realization space* of the matroid M . Points of this realization space correspond to isomorphism class of representations of M as hyperplane arrangements.

The thin Schubert cells are naturally torus torsors over realization spaces of hyperplane arrangements which we will now describe using the subtori $(\mathbf{k}^*)^{F_i}$ from the above subsection following [53], which summarizes the discussion in [60, Sect. 1.6].

Proposition 9.7 ([60, 1.6]). *If M is connected, then \mathbf{R}_M is the quotient of Gr_M by the free action of $(\mathbf{k}^*)^{n+1}/\mathbf{k}^*$. If M is not connected, there is a natural isomorphism*

$$\mathbf{R}_M \cong \mathbf{R}_{M|_{F_0}} \times \cdots \times \mathbf{R}_{M|_{F_k}}$$

where $M = M|_{F_0} \oplus \cdots \oplus M|_{F_k}$ is the direct sum decomposition. In this case, the $(\mathbf{k}^*)^{n+1}/(\mathbf{k}^*)$ -action naturally factors through the projection

$$(\mathbf{k}^*)^{n+1}/\mathbf{k}^* \rightarrow (\mathbf{k}^*)^{F_0}/\mathbf{k}^* \times \cdots \times (\mathbf{k}^*)^{F_k}/\mathbf{k}^*.$$

9.6 Matroid Representability and Universality

In a certain sense, the question of the representability of matroids is maximally complicated as a consequence of Mněv’s universality theorem which states that any possible singularity defined over \mathbb{Z} occurs on some thin Schubert cell (up to a natural equivalence). This has applications to moduli problems and logic. Here we state Mněv’s theorem in the form applied by Lafforgue [60].

Theorem 9.8. *Let X be an affine scheme of finite type over $\text{Spec } \mathbb{Z}$. Then there exists a connected rank 3 matroid M , an integer $N \geq 0$, and an open set $U \subseteq X \times \mathbb{A}^N$ projecting onto X such that U is isomorphic to the realization space \mathbf{R}_M .*

A very brief proof is given in [60]. More details in a slightly different context can be found in [62]. An accessible treatment that only considers the set-theoretic case is [81]. A new proof giving explicit bounds on the size of the matroid is given in [18]. The central idea of most proofs of this theorem, following Shor [88], is to add auxiliary variables and present X as being cut out by a very large system of equations each involving three variables and a single arithmetic operation. This new system is encoded as a line arrangement by using the von Staudt constructions (see [81] for an illustration).

This theorem quickly gives interesting counterexamples. If X is disconnected, then the realization space \mathbf{R}_M will be disconnected and the matroid will have representations (as hyperplane arrangements) that cannot be continuously deformed into each other. One can also control the fields over which X has a representation: a matroid M is representable over \mathbf{k} if and only \mathbf{R}_M has a \mathbf{k} -point. If $X = \text{Spec } \mathbb{F}_2$ and \mathbf{k} is some infinite field, then X and consequently \mathbf{R}_M only has \mathbf{k} -points if and only if $\text{char } \mathbf{k} = 2$. Similarly, if $X = \text{Spec } \mathbb{Z}[x]/(2x - 1)$, then \mathbf{R}_M only has \mathbf{k} -points if and only if $\text{char } \mathbf{k} \neq 2$.

As observed by Sturmfels [93], a decision procedure to determine if a matroid is representable over \mathbb{Q} is equivalent to a decision procedure to determine whether a system of polynomial equations has a solution in \mathbb{Q} . The existence of the decision procedure is a famous unsolved problem, Hilbert’s Tenth Problem over \mathbb{Q} . See Poonen’s survey [77] for more details about this problem.

For finite fields \mathbf{k} , the situation is a bit different. Because \mathbf{R}_M is isomorphic to an open set $U \subseteq X \times \mathbb{A}^N$, and it is possible for some \mathbf{k} -points to be missing from $U \subseteq X \times \mathbb{A}^N$, it is not necessarily true that the existence of \mathbf{k} -points in X is equivalent to the representability of M over \mathbf{k} . Moreover, the number of points of $X(\mathbf{k})$ and $\mathbf{R}_M(\mathbf{k})$ may be drastically different. For this reason, it's possible to have unique representability theorems similar to Theorem 6.3 for finite fields.

Mnëv's theorem is the primordial example of Murphy's Law in algebraic geometry. Murphy's Law, due to Vakil [99], is a theorem that states that certain classes of moduli spaces contain every singularity type over \mathbb{Z} up to a particular natural equivalence. In fact, Vakil was able to prove that a large number of moduli spaces obey Murphy's Law by relating them to realization spaces of matroids.

Matroid representability has been applied in tropical geometry to the study of linear systems on graphs in recent work of Cartwright [18], Jensen [50], and Len [63]. In particular, Mnëv's theorem has been used to construct linear systems on graphs that have pathological lifting behavior.

9.7 Matroid Representability and Semifields

One can view matroids as points of thin Schubert cells over field-like objects like blueprints or partial fields. This is the point of view taken in Dress's matroids over fuzzy rings [29] as studied by Dress with Wenzel [30]. Here, we use the notation of Lorscheid's blueprints and blue schemes [65] for the following highly speculative subsection.

Definition 9.9. A *blueprint* B is a multiplicatively written monoid A with neutral element 1 and absorbing element 0 together with an equivalence relation \mathcal{R} (given by \equiv) on the semiring $\mathbb{N}[A]$ of finite formal sums of elements of A such that

- (1) the relation \mathcal{R} is closed under addition and multiplication,
- (2) the absorbing element is equivalent to the empty sum, and
- (3) for $a, b \in A$ with $a \equiv b$, then $a = b$.

For example, one can define an object called the field of one element \mathbb{F}_1 to be the blueprint with $A = \{0, 1\}$ and \mathcal{R} to be the empty relation. Moreover, given any semiring R , one can produce a blueprint by setting A to be R^\times , where R^\times is R , considered as a multiplicative monoid, and setting

$$\mathcal{R} = \left\{ \sum a_i \equiv \sum b_i \mid \sum a_i = \sum b_i \text{ in } R \right\}.$$

In particular every field and idempotent semiring can be interpreted as a blueprint. An important blueprint is the none-some semifield \mathbb{B}_1 where we set $A = \{0, 1\}$ and $\mathcal{R} = \{1 + 1 \equiv 1\}$. Here we imagine a primitive culture with addition and multiplication but only numbers for 0 and more than 0. This is isomorphic to the min-plus semifield $\{0, \infty\}$ by taking $-\log$.

For a blueprint B , we can view a linear subspace defined over B as a B -point of $\text{Gr}(d + 1, n + 1)$ provided that we can treat $\text{Gr}(d + 1, n + 1)$ in its Plücker embedding as a model of $\text{Gr}(d + 1, n + 1)$ over B . While there are various rival notions for the definition of a B -point, one notion is an assignment of the Plücker coordinates to elements of A such that Plücker relations are satisfied. Under that definition, a matroid is a \mathbb{B}_1 -point. The value of a Plücker coordinate in \mathbb{B}_1 determines whether or not it is zero. The Plücker relations in \mathbb{B}_1 become the basis exchange axiom. In the case that B is induced from a field \mathbf{k} , then a B -point of $\text{Gr}(d + 1, n + 1)$ is the same thing as a \mathbf{k} -point.

Regular matroids should perhaps be thought of as \mathbb{F}_{12} -points of $\text{Gr}(d + 1, n + 1)$. Here \mathbb{F}_{12} , the degree 2 cyclotomic extension of \mathbb{F}_1 , is the blueprint generated by $A = \{0, 1, -1\}$ subject to the relation $1 + (-1) \equiv 0$. This is because regular matroids are those representable by totally unimodular matrices and therefore, all the Plücker coordinates are 0, 1, or -1 . An important idea in the study of the field of one element is the numerical prediction that the number \mathbb{F}_1 -points is the limit of the number of \mathbb{F}_q -points as $q \rightarrow 1$. Is there a similar prediction for the number of regular matroids?

Excluded minor characterizations of matroids can be thought of as matroid lifting results. Specifically, we have a morphism of blueprints $B' \rightarrow B$, and a B -point of the Grassmannian, and we ask when this B -point is in the image of the induced map $\text{Gr}(d + 1, n + 1)(B') \rightarrow \text{Gr}(d + 1, n + 1)(B)$. For example, there is a natural blueprint morphism $\mathbb{F}_{12} \rightarrow \mathbb{F}_2$. If M is a matroid that is representable over \mathbb{F}_2 , the question of whether it is regular corresponds to whether it lifts from a \mathbb{F}_2 -point to a \mathbb{F}_{12} -point. This is the point of view taken in the work of Pendavingh and van Zwam [75] phrased in the language of partial fields which include the usual fields and \mathbb{F}_{12} . They are able to give unified proofs of a number of representability results including the following theorem of Tutte:

Theorem 9.10. *A matroid is regular if and only if it is representable over both \mathbb{F}_2 and \mathbb{F}_3 .*

Here, they manipulate matrix representations of matroids by matrix operations, but it would be very interesting to rephrase their work in the language of deformation theory. The category of partial fields is more interesting than that of fields because of the existence of tensor products and many more morphisms.

9.8 *K-theoretic Matroid Invariants*

Speyer [90] introduced a way of producing valuative invariants of matroids by using the K -theory of the Grassmannian. He begins with an invariant valued in the K -theory of the Grassmannian. This invariant can be specialized by particular geometric operations to take values in less exotic rings. In fact, it is the main result of [36] that this invariant specializes to the Tutte polynomial.

We first review algebraic K -theory. See [37, Sect. 15.1] for a very brief summary. For a scheme X , let $K_0(X)$ be the Grothendieck group of coherent sheaves on X . It is the group of formal linear combinations of coherent sheaves on X subject to the relation

$$[\mathcal{F}] = [\mathcal{F}'] + [\mathcal{F}'']$$

whenever there is an exact sequence of coherent sheaves

$$0 \longrightarrow \mathcal{F}' \longrightarrow \mathcal{F} \longrightarrow \mathcal{F}'' \longrightarrow 0.$$

Given a proper morphism of schemes $f : X \rightarrow Y$, there is a pushforward homomorphism $f_* : K_0(X) \rightarrow K_0(Y)$ given by

$$f_* \mathcal{F} = \sum_{i=0}^{\infty} (-1)^i [R^i f_* \mathcal{F}]$$

On the other hand, we may define $K^0(X)$ as the Grothendieck group of vector bundles on X , the group of formal linear combinations of vector bundles on X subject to the relation

$$[E] = [E'] + [E'']$$

whenever E' is a sub-bundle of E with quotient $E'' = E/E'$. Given a morphism $f : X \rightarrow Y$, there is a pullback homomorphism $f^* : K^0(Y) \rightarrow K^0(X)$ given by pullback of vector bundles. The group $K^0(X)$ can be made into a ring by introducing tensor product as multiplication. The natural group homomorphism $K^0(X) \rightarrow K_0(X)$ taking a vector bundle to its sheaf of sections is an isomorphism if X is a non-singular variety. We will freely switch between $K_0(X)$ and $K^0(X)$ as all of our schemes will be non-singular varieties. We will make use of the fact that $K_0(\mathbb{P}^n) = \mathbb{Q}[x]/(x^{n+1})$ where $x = 1 - \mathcal{O}(-1)$ which corresponds to the structure sheaf of a hyperplane.

Let $T = (\mathbf{k}^*)^{n+1}$ be an algebraic torus acting on a scheme X . We may define equivariant K -groups $K_0^T(X)$ and $K_T^0(X)$ by considering equivariant coherent sheaves and vector bundles. There are non-equivariant restriction maps

$$K_0^T(X) \rightarrow K_0(X), \quad K_T^0(X) \rightarrow K^0(X)$$

forgetting the T -action. Now, the equivariant K -theory of a point is particularly simple: $K_T^0(\text{pt})$ is the Grothendieck group of representations of T . By taking characters, we see $K_T^0 = \mathbb{Z}[t_0^{\pm}, \dots, t_n^{\pm}]$, the ring of Laurent polynomials.

The Grassmannian has a natural group action which allows us to consider equivariant K -theory [58]. By localization, this equivariant K -theory has a quite combinatorial flavor. Let T act on $\text{Gr}(d + 1, n + 1)$ as follows: $t \in T$ dilates the

coordinates of ambient \mathbf{k}^{n+1} taking a subspace $P \in \text{Gr}(d + 1, n + 1)$ to $t \cdot P$. Note that the diagonal subtorus in T acts trivially. The equivariant K -theory of the Grassmannian can be expressed in terms of its fixed points and 1-dimensional orbits under T . The following is easily verified:

Lemma 9.11. *The fixed points of the T -action on $\text{Gr}(d + 1, n + 1)$ are in one-to-one correspondence with subsets $B \subset \{0, 1, \dots, n\}$ with $|B| = d + 1$ with the fixed points x_B given by the $(d + 1)$ -dimensional subspaces of the form $\text{Span}(e_i \mid i \in B)$. The 1-dimensional torus orbits of the action are in one-to-one correspondence with pairs $B_1, B_2 \subset \{0, 1, \dots, n\}$ with $|B_1| = |B_2| = d + 1$ and $|B \Delta B'| = 2$ with the closure of the torus orbit parameterized by \mathbb{P}^1 as*

$$[s : t] \mapsto \text{Span}(e'_i \mid i' \in B_1 \cap B_2) + \mathbf{k}(se_i + te_j)$$

where $i \in B_1 \setminus B_2$ and $j \in B_2 \setminus B_1$.

Because $\text{Gr}(d + 1, n + 1)$ is a smooth projective variety, has finitely many fixed points, and the closure of each 1-dimensional torus orbits is isomorphic to \mathbb{P}^1 , $\text{Gr}(d + 1, n + 1)$ is equivariantly formal [101, Corollary 5.12], [58, Corollary A.5] which means that its equivariant K -theory is controlled by the fixed points and 1-dimensional orbits. Specifically, the restriction map to the fixed points

$$K_T^0(\text{Gr}(d + 1, n + 1)) \rightarrow K_T^0(\text{Gr}(d + 1, n + 1)^T) = \bigoplus_B K_T^0(x_B)$$

is an injection whose image is given by $f_B \in K_0^T(x_B)$ for all B satisfying

$$f_{S \cup \{i\}} \equiv f_{S \cup \{j\}} \pmod{1 - t_i/t_j}$$

where $S \in \binom{\{0, \dots, n\}}{d}$, $i, j \notin S$.

Now, we describe Speyer’s matroid invariant. Let $P \in \text{Gr}(d + 1, n + 1)$. The invariant is produced from the K -theory class of the orbit closure $\overline{T \cdot P}$,

$$\mathcal{O}_{\overline{T \cdot P}} \in K_0(\text{Gr}(d + 1, n + 1)).$$

Because $\overline{T \cdot P}$ is T -invariant, its structure sheaf is T -equivariant. The corresponding equivariant K -theory class, $y(P) \in K_0^T(\text{Gr}(d + 1, n + 1))$ can be given in terms of the matroid polytope. We will give the description of its restrictions $y(P)_B \in K_0^T(x_B)$. For B , a basis of the matroid, let $\text{Cone}_B(P(M))$ be the cone of $P(M)$ at B , in other words the real nonnegative span of all vectors of the form $v - e_B$ for $v \in P(M)$. The Hilbert series of a cone C is given by

$$\text{Hilb}(C) = \sum_{a \in \mathbb{C} \cap \mathbb{Z}^{n+1}} t^a.$$

We define $y(P)$ by

$$y(P)_B = \begin{cases} \text{Hilb}(\text{Cone}_B(P(M))) \prod_{i \in B} \prod_{j \notin B} (1 - t_i^{-1} t_j) & \text{if } B \text{ is a basis of } P \\ 0 & \text{else} \end{cases}$$

The rational function $y(P)_B$ turns out to be a Laurent polynomial. Note that $y(P)$ depends only on the matroid of P and so can be written $y(M)$. This definition gives a well-defined class of $K_0^T(\text{Gr}(d + 1, n + 1))$ even when M is not representable. This invariant turns out to be valuative although it is not universal among valuative invariants.

The invariant $y(M)$ specializes to the Tutte polynomial in a geometric fashion. Recall that $\text{Fl}(d_1, \dots, d_r; n + 1)$ is the flag variety whose points parameterize flags of projective subspaces $F_{d_1-1} \subset \dots \subset F_{d_r-1} \subset \mathbb{P}^n$ where $\dim F_i = i - 1$. There is a natural inclusion

$$\text{Fl}(1, n; n + 1) \hookrightarrow \mathbb{P}^n \times \mathbb{P}^n$$

taking $F_1 \subset F_{n-1}$ to (F_1, F_{n-1}) where the two \mathbb{P}^n 's parameterize points and hyperplanes in \mathbb{P}^n . There is a diagram (borrowed from [36]) of morphisms forgetting various stages of the flags:

$$\begin{array}{ccc}
 & \text{Fl}(1, d + 1, n; n + 1) & \\
 \swarrow \pi_{d+1} & & \searrow \\
 \text{Gr}(d + 1, n + 1) & & \text{Fl}(1, n; n + 1) \\
 & & \searrow \pi_{1n} \\
 & & \mathbb{P}^n \times \mathbb{P}^n
 \end{array} \tag{2}$$

If we write $K^0(\mathbb{P}^n \times \mathbb{P}^n) = \mathbb{Q}[x, y]/(x^{n+1}, y^{n+1})$, one has

Theorem 9.12 ([36]). *Let H be the pullback of $\mathcal{O}(1)$ on $\mathbb{P}^{\binom{n+1}{d+1}-1}$ to $\text{Gr}(d + 1, n + 1)$ by the Plücker embedding. Interpreting $y(M)$ as a non-equivariant class, we have*

$$(\pi_{1n})_* \pi_{d+1}^*(y(M) \cdot [H]) = T_M(x, y).$$

The proof works by relating the K -theoretic class to the rank-generating polynomial.

10 Review of Toric Varieties

We give a quick review of toric varieties in preparation for the next section. More complete references are [19, 38].

10.1 Rational Fans and Minkowski Weights

A toric variety $X = X(\Delta)$ is an algebraic variety specified by a rational fan Δ in $N_{\mathbb{R}} = N \otimes_{\mathbb{Z}} \mathbb{R}$ for a lattice $N \cong \mathbb{Z}^n$. They are normal varieties compactifying the algebraic torus $T = (\mathbf{k}^*)^n$ such that the natural multiplication of T on itself extends to a T -action. A rational fan is a particular set, Δ of strongly convex rational polyhedral cones. A strongly convex rational polyhedral cone is the nonnegative span of finitely many vectors in N which contains no lines through the origin. A rational fan is defined as a polyhedral complex whose cells are strongly convex rational polyhedral cones, that is:

1. if $\sigma \in \Delta$, then every face of σ is an element of Δ , and
2. if $\sigma, \sigma' \in \Delta$, then $\sigma \cap \sigma' \in \Delta$.

We say that a fan is pure of dimension d if every maximal cone is d -dimensional. For a cone σ , let σ° denote the relative interior of σ . Associated with a fan Δ is a toric variety $X(\Delta)$ defined over a field \mathbf{k} .

The toric variety is stratified by torus orbits. Write N^\vee for the lattice dual to N . Write N_σ for the sublattice of N given by $\text{Span}(\sigma) \cap N$. For a cone σ , let $\sigma^\perp \subseteq N^\vee$ be the kernel of the projection $N^\vee \rightarrow N_\sigma^\vee$ that is dual to the natural inclusion. Given a cone σ with $\dim(\sigma) = k$, there is a torus orbit \mathcal{O}_σ with $\dim(\mathcal{O}_\sigma) = n - k$. This torus orbit is canonically isomorphic to $\text{Hom}(\sigma^\perp, \mathbf{k}^*)$. The torus orbit $\mathcal{O}_0 = \text{Hom}(N^\vee, \mathbf{k}^*)$ corresponding to $0 \in \Delta$ is called the *big open torus* and is isomorphic to T . We can view elements of N^\vee as regular functions on T . Specifically, we view $m \in N^\vee$ as the function χ^m , for $t \in \text{Hom}(N^\vee, \mathbf{k}^*)$,

$$\chi^m : t \mapsto t(m).$$

We refer to χ^m as a character of T . The group T acts on \mathcal{O}_σ as follows: for $\gamma : \sigma^\perp \rightarrow \mathbf{k}^*$, we have $t \cdot \gamma$ given by for $m \in \sigma^\perp$,

$$t \cdot \gamma(m) = \chi^m(t)\gamma(m).$$

The toric variety $X(\Delta)$ can be decomposed (as a set with action of T as)

$$X(\Delta) = \bigsqcup_{\sigma \in \Delta} \mathcal{O}_\sigma.$$

The closure of the orbit \mathcal{O}_σ is denoted by $V(\sigma)$. The cones of Δ are in inclusion-reversing bijection with orbit closures of $X(\Delta)$: if τ is a face of σ in Δ , then $V(\sigma) \subseteq V(\tau)$. In other words, $\mathcal{O}_\sigma \subseteq \overline{\mathcal{O}_\tau}$ if and only if $\tau \subseteq \sigma$. The closure of $T = \mathcal{O}_0$ is $X(\Delta)$. Characters of T extend to $X(\Delta)$ as rational functions.

A toric variety $X(\Delta)$ is compact if and only if the *support of Δ* , $\bigcup \sigma$ is equal to $N_{\mathbb{R}}$. A toric variety is smooth if and only if every cone is unimodular, that is, it is the nonnegative span of a set of vectors that can be extended to a basis of the lattice N .

Example 10.1. The building blocks of smooth toric varieties are the toric varieties associated with unimodular simplices. Let $N = \mathbb{Z}^n$. Write the basis vectors of N as e_1, \dots, e_n . Let σ be the nonnegative span of e_1, \dots, e_k . Let Δ consist of σ and its faces. Then $X(\Delta) = \mathbf{k}^k \times (\mathbf{k}^*)^{n-k}$. Write the coordinates on $X(\Delta)$ as $(x_1, \dots, x_k, x_{k+1}, \dots, x_n)$. The faces of σ are

$$\sigma_S = \text{Span}_{\geq 0}(e_i \mid i \in S)$$

for $S \subseteq \{1, \dots, k\}$. The torus orbits and their closures are given by

$$\begin{aligned} \mathcal{O}_{\sigma_S} &= \{x \mid x_i = 0 \text{ if and only if } i \in S\}, \\ V(\sigma_S) &= \{x \mid x_i = 0 \text{ if } i \in S\}. \end{aligned}$$

The group $(\mathbf{k}^*)^n$ acts by multiplication on coordinates. The characters are the usual monomials in the x_i 's.

Example 10.2. The most basic compact toric variety is \mathbb{P}^n . Let $N = \mathbb{Z}^{n+1}/\mathbb{Z}$ where \mathbb{Z}^{n+1} is spanned by unit vectors e_0, e_1, \dots, e_n and the quotient is by the span of the diagonal element $e_0 + e_1 + \dots + e_n$. The cones of $\Delta_{\mathbb{P}^n}$ are given by $\sigma_S = \text{Span}_{\geq 0}(e_i \mid i \in S)$ for each proper subset $S \subsetneq \{0, 1, \dots, n\}$. Then $X(\Delta_{\mathbb{P}^n}) = \mathbb{P}^n$. The torus orbits and their closures are given in homogeneous coordinates by

$$\begin{aligned} \mathcal{O}_{\sigma_S} &= \{[X_0 : X_1 : \dots : X_n] \in \mathbb{P}^n \mid X_i = 0 \text{ if and only if } i \in S\}, \\ V(\sigma_S) &= \{[X_0 : X_1 : \dots : X_n] \in \mathbb{P}^n \mid X_i = 0 \text{ if } i \in S\}. \end{aligned}$$

The group $T = (\mathbf{k}^*)^{n+1}/\mathbf{k}^*$ acts on \mathbb{P}^n by

$$(t_0, t_1, t_2, \dots, t_n) \cdot [X_0 : X_1 : \dots : X_n] = [t_0 X_0 : t_1 X_1 : \dots : t_n X_n].$$

Note that the diagonal of $(\mathbf{k}^*)^{n+1}$ acts trivially. The characters are rational functions on \mathbb{P}^n : N^\vee is canonically isomorphic to the sublattice of vectors $(m_0, m_1, \dots, m_n) \in \mathbb{Z}^{n+1}$ satisfying $m_0 + m_1 + \dots + m_n = 0$, and

$$\chi^m([X_0 : X_1 : \dots : X_n]) = \prod_i X_i^{m_i}.$$

There is an important notion of morphisms of toric varieties. Let N and N' be lattices with rational fans Δ, Δ' in $N_{\mathbb{R}}, N'_{\mathbb{R}}$, respectively. Let $\phi : N' \rightarrow N$ be a homomorphism of lattices such that for any cone $\sigma' \in \Delta'$, there is a cone $\sigma \in \Delta$ with $\phi_{\mathbb{R}}(\sigma') \subset \sigma$. In this case, we say $\phi : \Delta' \rightarrow \Delta$ is a morphism of fans. Then, there is an induced map of toric varieties $\phi_* : X(\Delta') \rightarrow X(\Delta)$ that intertwines the torus actions. Of particular interest for us is the case where $N = N'$ and the homomorphism ϕ is the identity. In that case, Δ' is a *refinement* of Δ , that is, every cone $\sigma' \in \Delta'$ is contained in a cone of Δ . The induced morphism $\phi_* : X(\Delta') \rightarrow X(\Delta)$ is a birational morphism.

Example 10.3. Blow-ups of $X(\Delta)$ at the orbit closures $V(\sigma)$ can be phrased in terms of refinements of fans [38, Sect. 2.4], [19, Definition 3.3.17]. We will explain the case of a unimodular cone as in Example 10.1. The general case is similar. Let σ be the cone spanned by e_1, \dots, e_n . Let Δ be the fan consisting of all faces of σ . Let e' be the barycenter of σ given by $e' = e_1 + \dots + e_n$. Let Δ' be the fan consisting of all cones spanned by subsets of $\{e', e_1, \dots, e_n\}$ not containing $\{e_1, \dots, e_n\}$. By an explicit computation, one can show that $X(\Delta')$ is the blow-up of \mathbf{k}^n at the origin. This has the effect of subdividing the cone σ while not changing its boundary.

We can apply this operation to a top-dimensional cone of a toric variety. Let Δ be a fan, and let σ be a top-dimensional, unimodular cone of Δ . We can form $\tilde{\Delta}$ by replacing σ with the cones considered above and not changing any of the other cones. Then $X(\tilde{\Delta})$ is the blow-up of $X(\Delta)$ at the smooth point $V(\sigma)$.

This operation can also be applied to smaller-dimensional cones as well. Here, we follow the exposition of [19]. Let $X(\Delta)$ be a smooth toric variety. For a cone σ , let $\sigma(1)$ be the set of primitive lattice vectors through the 1-dimensional faces of σ . Recall that a primitive lattice vector is a vector $v \in N$ such that if $v = nw$ for $n \in \mathbb{Z}$ and $w \in N$ then $n = \pm 1$. Let τ be a cone of Δ . Let the barycenter of τ be

$$u_{\tau} = \sum_{u_{\rho} \in \tau(1)} u_{\rho}.$$

For each cone σ containing τ , define a fan

$$\Delta'_{\sigma}(\tau) = \{\text{Span}_{\geq 0}(S) \mid S \subseteq \{u_{\tau}\} \cup \sigma(1), \tau(1) \not\subseteq S\}.$$

Then the subdivision of Δ relative to τ is the fan

$$\Delta'(\tau) = \{\sigma \in \Delta \mid \sigma \not\supseteq \tau\} \cup \bigcup_{\sigma \supseteq \tau} \Delta'_{\sigma}(\tau).$$

It turns out that $X(\Delta'(\tau))$ is the blow-up of $X(\Delta)$ along the subvariety $V(\tau)$. The orbit closure $V(\mathbb{R}_{\geq 0}u_{\tau})$ is the exceptional divisor of the blow-up. Observe that only the cones containing τ are affected by the subdivision.

We can form the barycentric subdivision of Δ by first subdividing the cones τ with $\dim(\tau) = n$ and then subdividing the cones τ of Δ with $\dim(\tau) = n - 1$ and

so on down to the 2-dimensional cones. This produces an iterated blow-up of $X(\Delta)$. This construction will be very important in the sequel.

A natural way that rational fans arise is as normal fans to lattice polytopes. Let P be a lattice polytope in \mathbb{R}^n , that is, a polytope whose vertices are in \mathbb{Z}^n . Set $N = \text{Hom}(\mathbb{Z}^n, \mathbb{Z})$.

Definition 10.4. For a face Q of P , the (*inward*) *normal cone* to Q is the set

$$\sigma_Q = \{w \in N \mid P_w \supseteq Q\}.$$

For w in the relative interior of σ_Q , we have $P_w = Q$.

Definition 10.5. The *inward normal fan*, $N(P)$ of the polytope P is the union of the cones σ_Q as Q ranges over the faces of P .

The correspondence between Q and σ_Q is inclusion-reversing. If P is full-dimensional, then $N(P)$ is strongly convex. Because P is an integral polytope, $N(P)$ is a rational fan.

We can use normal fans to relate not-necessarily-normal toric varieties to the more usual toric varieties. Suppose we have $i : \mathbb{P}_{\mathcal{A}} \rightarrow \mathbb{P}^N$. If i is an inclusion and $\mathbb{P}_{\mathcal{A}}$ is normal, then $\mathbb{P}_{\mathcal{A}}$ is the toric variety associated with the normal fan of the weight polytope of $\mathbb{P}_{\mathcal{A}}$.

The fan $\Delta_{\mathbb{P}^n}$ occurs as the normal fan to the simplex whose vertices are $-e_0, -e_1, \dots, -e_n$ in the hyperplane defined by $x_0 + x_1 + \dots + x_n = -1$ in \mathbb{R}^{n+1} .

10.2 Intersection Theory and Minkowski Weights

We review the central notions of intersection theory [37], eventually specializing to toric varieties. We recommend [38] as a reference for toric varieties, [39] for results on intersection theory on toric varieties, and [51] for results most directly suited to our purposes.

Let X be an n -dimensional algebraic variety over \mathbf{k} . The *cycle group* $Z_p(X)$ is the group of finite formal integer combinations $\sum n_i[Z_i]$ where each Z_i is an irreducible p -dimensional subvariety of X . These sums are called cycles. If all the coefficients are nonnegative, the cycle is said to be effective. A cycle is declared to be rationally equivalent to 0 if there exists irreducible $(p + 1)$ -dimensional varieties W_1, \dots, W_l together with rational functions f_1, \dots, f_l on W_1, \dots, W_l , respectively, such that

$$\sum n_i Z_i = \sum_j (f_j)$$

where (f_j) denotes the principal divisor on W_j associated with the rational function f_j . The Chow group $A_p(X)$ is the quotient of $Z_p(X)$ by the subgroup of cycles rationally equivalent to 0. Elements of $A_p(X)$ are cycle classes, but we may refer to them

as cycles when we have picked a member of their class. A cycle class is said to be *effective* if it is rationally equivalent to an effective cycle. The Chow groups should be thought of as an algebraic analogue of the homology groups $H_{2p}(X)$ with the caveat that not every homology class is realizable by a cycle and that rational equivalence is a much finer relation than homological equivalence.

If X is smooth, then there is an intersection product

$$\cdot : A_p(X) \otimes A_{p'}(X) \rightarrow A_{p+p'-n}(X)$$

which should be thought of as taking two irreducible subvarieties Z, Z' to the rational equivalence class of their intersection $[Z \cap Z']$ if the subvarieties meet transversely. A lot of work has to be done to make this intersection product well defined for pairs of subvarieties that do not meet transversely. The intersection product is well defined on rational equivalence classes. We think of elements of $A_0(X)$ as a formal sum of points, and if X is compact, there is a degree map:

$$\text{deg} : A_0(X) \rightarrow \mathbb{Z}$$

taking $\sum n_i P_i$ to $\sum n_i$. When X is smooth and compact, we may define the Chow cohomology groups simply as $A^k(X) = A_{n-k}(X)$. These groups are, in fact, graded rings under the cup product

$$\cup : A^k(X) \otimes A^{k'}(X) \rightarrow A^{k+k'}(X)$$

which is simply the intersection product. We will make use of the cap-product

$$\cap : A^k(X) \otimes A_p(X) \rightarrow A_{p-k}(X)$$

where $c \cap z$ is given by taking the intersection product of c (considered as a $(n - k)$ -dimensional cycle) with z considered as a p -dimensional cycle. By definition, for $c, d \in A^*(X), z \in A_*(X)$,

$$c \cap (d \cap z) = (c \cup d) \cap z.$$

For non-smooth varieties, the definition of Chow cohomology is quite different. Here, we used Poincaré duality to simplify our exposition.

Definition 10.6. A Chow cohomology class $c \in A^1(X(\Delta))$ is said to be numerically effective or *nef* if for every curve C on X , $\text{deg}(c \cap [C]) \geq 0$.

For $X(\Delta)$, a compact toric variety, the Chow cohomology groups $A^k(X(\Delta))$ are canonically isomorphic to a combinatorially defined object, the group of Minkowski weights. Let $\Delta^{(k)}$ denote the set of all cones in Δ of dimension $n - k$. If $\tau \in \Delta^{(k+1)}$ is contained in a cone $\sigma \in \Delta^{(k)}$, let $v_{\sigma/\tau} \in N/N_\tau$ be the primitive generator of the ray $(\sigma + N_\tau)/N_\tau$.

Definition 10.7. A function $c : \Delta^{(k)} \rightarrow \mathbb{Z}$ is said to be a *Minkowski weight* of codimension k if it satisfies the *balancing condition*, that is, for every $\tau \in \Delta^{(k+1)}$,

$$\sum_{\substack{\sigma \in \Delta^{(k)} \\ \sigma \supset \tau}} c(\sigma)v_{\sigma/\tau} = 0$$

in N/N_τ . The *support* of c is the set of cones in $\Delta^{(k)}$ on which c is non-zero.

Theorem 10.8 ([39]). *The Chow group $A^k(X(\Delta))$ is canonically isomorphic to the group of codimension k Minkowski weights.*

The correspondence between Chow cohomology classes and Minkowski weights is as follows: given $d \in A^k(X)$, define $c(\sigma) = \deg(d \cap [V(\sigma)])$. The content of the above theorem is that Chow cohomology classes are determined by their intersections with orbit closures. The balancing condition is a combinatorial translation of the fact that cohomology classes are constant on rational equivalence classes generated by the rational functions given by characters of the torus \mathcal{O}_τ on the orbit closure $V(\tau)$. The degree $\deg(c)$ of a class $c \in A^n(X)$ is defined to be $c(0)$, the value of c on the unique zero-dimensional cone 0 . There is a combinatorial description of the cup product of Chow cohomology classes by the fan-displacement rule which we will not describe. The identity in the Chow ring is $c \in A^0(X(\Delta))$ given by

$$c(\sigma) = 1 \text{ for } \sigma \in \Delta^{(0)}.$$

For compact smooth toric varieties, Chow cohomology is isomorphic to singular cohomology.

Example 10.9. Let us consider $\Delta_{\mathbb{P}^n}$. We know that $A^k(\mathbb{P}^n) = \mathbb{Z}H^k$ where H is a hyperplane class, and, consequently, H^k is the class of a codimension k projective subspace. Now, we will see this fact in terms of Minkowski weights. Let c be a Minkowski weight of codimension k . It is straightforward to verify that the Minkowski weight condition is equivalent to c being constant on $\Delta^{(k)}$. If $\tau \in \Delta^{(k)}$, $V(\tau)$ is a k -dimensional coordinate subspace. Now, c , considered as a cycle, must intersect $V(\tau)$ in $c(\tau)$ points counted with multiplicity. Therefore, c is in the class of $c(\tau)H^k$.

There is an equivariant version of Chow cohomology which has a combinatorial description on toric varieties. Here, equivariant Chow cohomology means equivariant with respect to the T -action and is analogous to equivariant cohomology. We will only make use of the codimension 1 case. Let $N^\vee = \text{Hom}(N, \mathbb{Z})$ be the dual lattice to N , interpreted as linear functions on $N_{\mathbb{R}}$ with integer slopes. A linear function on a cone σ can be interpreted as an element of $N^\vee(\sigma) = N^\vee/\sigma^\perp$. A piecewise linear function α on Δ is a continuous function on the support of Δ whose restriction to each cone is a linear function with integer slopes, or more formally as the following:

Definition 10.10. A piecewise linear function on the fan Δ is a collection of elements $\alpha_\sigma \in N^\vee/\sigma^\perp$ for each $\sigma \in \Delta$ such that for $\tau \subset \sigma$, the image of α_σ under the quotient $N^\vee/\sigma^\perp \rightarrow N^\vee/\tau^\perp$ is α_τ .

Observe that if $\phi : \Delta' \rightarrow \Delta$ is a morphism of fans, and α is a piecewise linear function on Δ , then the pullback $\phi^*\alpha$ is a piecewise linear function of Δ' . The piecewise linear function α can be thought of as a T -Cartier divisor on $X(\Delta)$: if the restriction of α to σ , α_σ is given by $m \in M/M(\sigma)$, and the Cartier divisor is locally defined by the rational function χ^m , the character associated with m on the toric open affine associated with σ . This T -Cartier divisor can be thought of as an element of $A_T^1(X(\Delta))$, the equivariant Chow cohomology group [31]. There is a natural non-equivariant restriction map $i^* : A_T^1(X(\Delta)) \rightarrow A^1(X(\Delta))$ that takes the T -Cartier divisor to an ordinary Chow cohomology class. The piecewise linear function α induces a T -equivariant line bundle L on $X(\Delta)$ and $i^*\alpha = c_1(L)$, the first Chern class of L .

If $c \in A^k(X(\Delta))$ is viewed as a Minkowski weight, we may compute the cup product $i^*\alpha \cup c$ as an element of $A^{k+1}(X)$ by using a formula that first appeared in [5]: for $\sigma \in \Delta^{(k)}$, $\tau \in \Delta^{(k+1)}$, let $u_{\sigma/\tau}$ be a vector in N_σ descending to $v_{\sigma/\tau}$ in N/N_τ ; then the value of $i^*\alpha \cup c$ on a cone $\tau \in \Delta^{(k+1)}$ is

$$(i^*\alpha \cup c)(\tau) = - \sum_{\sigma \in \Delta^{(k)} | \sigma \supset \tau} \alpha_\sigma(u_{\sigma/\tau})c(\sigma) + \alpha_\tau \left(\sum_{\sigma \in \Delta^{(k)} | \sigma \supset \tau} c(\sigma)u_{\sigma/\tau} \right)$$

where α_σ (respectively α_τ) is the linear function on N_σ (on N_τ) which equals α on σ (on τ). Taking $i^*\alpha \cup c$ yields a Minkowski weight. The following lemma can be proved using intersection theory [52] or by elementary means [5, Proposition 3.7]:

Lemma 10.11. *Let α be a piecewise linear function on Δ , and let c be a Minkowski weight of codimension k . Then $i^*\alpha \cup c$ is a Minkowski weight of codimension $k + 1$.*

A T -Cartier divisor α is said to be nef if for every codimension 1 cone $\tau \in \Delta^{(1)}$, we have $i^*\alpha(\tau) \geq 0$. This says that the cohomology class $i^*\alpha$ is nonnegative on any 1-dimensional orbit closure. This happens if and only if $i^*\alpha$ is nonnegative on every curve class and is therefore nef in the classical sense. Consequently if α is nef, it induces a T -equivariant line bundle on $X(\Delta)$ whose first Chern class is nef. Nefness can be interpreted in terms of convexity. If τ is a codimension 1 cone of Δ , it is contained in two top-dimensional cones, σ_1, σ_2 . We can pick $u_{\sigma_1/\tau}, u_{\sigma_2/\tau}$ such that $u_{\sigma_1/\tau} + u_{\sigma_2/\tau} = 0$. Then, we have

$$i^*\alpha(\tau) = -\alpha_{\sigma_1}(u_{\sigma_1/\tau}) - \alpha_{\sigma_2}(u_{\sigma_2/\tau})$$

which expresses how the slope of α changes as we pass through the wall τ from σ_1 to σ_2 . Therefore, the condition $i^*\alpha(\tau) \geq 0$ can be interpreted as a convexity condition on the piecewise linear function α .

Example 10.12. Consider the fan $\Delta_{\mathbb{P}^n}$. Let w_1, \dots, w_n be coordinates on $N_{\mathbb{R}}$ given by the basis vectors e_1, \dots, e_n . Set $e_0 = -e_1 - \dots - e_n$. Let α be the piecewise linear function given by

$$\alpha = \min(0, w_1, \dots, w_n).$$

Let us compute $\iota^* \alpha \in A^{n-1}(\mathbb{P}^n)$. The $(n - 1)$ -dimensional cones of $\Delta_{\mathbb{P}^n}$ are of the following form: for $\{j, j'\} \subset \{0, 1, \dots, n\}$

$$\tau_{jj'} = \text{Span}(e_i \mid i \in \{0, 1, \dots, n\}, i \neq j, i \neq j'),$$

By explicit computation,

$$\iota^* \alpha(\tau_{jj'}) = 1.$$

The orbit closures $V(\tau_{jj'})$ are the coordinate lines in \mathbb{P}^n cut out by setting all but two homogeneous coordinates to 0. This tells us that $\iota^* \alpha$ intersects each 1-dimensional coordinate line in a point, and therefore, it is a hyperplane class.

If $\phi : \Delta' \rightarrow \Delta$ is a refinement of fans, then we can treat a piecewise linear function α on Δ as a piecewise linear function on Δ' , the pullback $\phi^* \alpha$. Moreover, if α is nef on Δ , then $\phi^* \alpha$ is nef on Δ' .

There is a way of associating a Minkowski weight of codimension r with a codimension r subvariety Y of a smooth, complete toric variety $X(\Delta)$. This Minkowski weight should be thought of as a Poincaré dual to the cycle $[Y]$. Define a function

$$c : \Delta^{(r)} \rightarrow \mathbb{Z}, \quad \sigma \mapsto \deg([Y] \cdot [V(\sigma)]).$$

Then c is a Minkowski weight, called the *associated cocycle* of Y . See [51] or [95] for details.

Definition 10.13. A d -dimensional irreducible subvariety Y of $X(\Delta)$ is said to *intersect the torus orbits of $X(\Delta)$ properly*, if for any cone σ of Δ ,

$$\dim(Y \cap V(\sigma)) = \dim(Y) + \dim V(\sigma) - n.$$

Lemma 10.14 ([51, Lemma 9.2]). *If c is the associated cocycle of a subvariety Y , then*

$$c \cap [X] = [Y] \in A_r(X).$$

If Y intersects orbits properly, the associated cocycle is essentially the same thing as the tropicalization of $Y^\circ \subset (\mathbf{k}^*)^n$ where Y° is defined below.

Definition 10.15. For a subvariety $Y \subset X(\Delta)$, the *interior* of Y is

$$Y^\circ = Y \cap (\mathbf{k}^*)^n$$

where $(\mathbf{k}^*)^n$ is the big open torus of $X(\Delta)$.

It is not difficult to show that the Chow ring of projective space, $A^*(\mathbb{P}^n)$ is $\mathbb{Z}[H]/H^{n+1}$ where H is the class of a hyperplane. Because the intersection of n hyperplanes in generic position is a point, $\deg(H^n) = 1$. Moreover, we have

$$A^*(\mathbb{P}^n \times \mathbb{P}^n) \cong \mathbb{Z}[H_1, H_2]/(H_1^{n+1}, H_2^{n+1})$$

where the codimension of a cohomology class is the total degree of the corresponding polynomial in H_1, H_2 . An element of $A^k(\mathbb{P}^n \times \mathbb{P}^n)$ can be written as

$$c = \sum_{i=0}^k a_i H_1^i H_2^{k-i}$$

for $a_i \in \mathbb{Z}$. Because $\deg(H_1^n H_2^n) = 1$, $a_i = \deg(H_1^{n-i} H_2^{n-k+i} \cup c)$.

11 Bergman Fans

11.1 Definition of a Bergman Fan

The Bergman fan is a combinatorial object associated with a matroid. We will see that it leads to a cryptomorphic definition of a matroid. It was first introduced as the logarithmic limit set by Bergman [8] and then shown to be a finite polyhedral complex by Bieri–Groves [9]. Sturmfels gave a combinatorial definition [94] which was elaborated by Ardila–Klivans [3].

The logarithmic limit set is an invariant of a $(d + 1)$ -dimensional linear subspace $V \subset \mathbb{C}^{n+1}$. The *amoeba* $A(V)$ of V is the set of all vectors of the form

$$(\log |x_0|, \log |x_2|, \dots, \log |x_n|) \in \mathbb{R}^{n+1}$$

for $x \in V$. One considers the Hausdorff limit of dilates $tA(V)$ as t goes to 0. This gives a logarithmic limit set, which is also called the tropicalization and captures the asymptotic behavior of the amoeba. Because V is invariant under dilation by elements of \mathbb{C}^* , the amoeba and hence the logarithmic limit set is invariant under translation by the diagonal vector $(1, \dots, 1)$. Therefore, we may view the logarithmic limit set as a subset of $\mathbb{R}^{n+1}/\mathbb{R}$. We will give a definition of the logarithmic limit set with reference only to the matroid of V and which makes sense for non-representable matroids.

Let M be a matroid on $E = \{0, 1, \dots, n\}$. We first define the underlying set of the Bergman fan. Recall that for $w \in \mathbb{R}^E$, by Definition 8.10, M_w is the matroid whose bases are the bases of M of minimal w -weight.

Definition 11.1. An element $w \in \mathbb{R}^E$ is said to be *valid* if M_w has no loops. The *underlying set of the Bergman fan* of M is the subset of \mathbb{R}^E consisting of valid w .

Observe that the underlying set of the Bergman fan is closed. If we have a sequence $\{w_j\}$ in the underlying set of the Bergman fan with $\lim w_j = w$, then $P(M)_w$ contains $P(M)_{w_j}$ for sufficiently large j . If M_{w_j} has no loops, every element of the ground set occurs as an element of some basis of M_{w_j} . To show that this is true for M_w , let $i \in E$. Then, for all j , $P(M)_{w_j}$ is not contained in the hyperplane $x_i = 0$. It follows that the same is true for M_w . Consequently, i is an element of some basis of $P(M)_w$ and hence is not a loop.

Proposition 11.2 ([94]). *The underlying set of the Bergman fan is the logarithmic limit set of any realization of M as a linear subspace in \mathbf{k}^{n+1} .*

Now, because the underlying set of the Bergman fan is invariant under translation by the vector $(1, 1, \dots, 1)$, we quotient by that vector. Consider the lattice $N = \mathbb{Z}^{n+1}/\mathbb{Z}(1, 1, \dots, 1)$. We write $N_{\mathbb{R}}$ for $\mathbb{R}^{n+1}/\mathbb{R}(1, 1, \dots, 1)$. We will treat N as dual to the lattice in the hyperplane $x_0 + x_1 + \dots + x_n = d + 1$ that contains the matroid polytope $P(M)$. The underlying set of the Bergman fan is the union of cones of the normal fan to the matroid polytope because the matroid M_w only depends on which cone of the normal fan to which w belongs. Therefore, the underlying set of the Bergman fan can be given the structure of a subfan of the normal fan of the matroid polytope. However, following [3], we will put a finer structure on the Bergman fan. This structure will be very important in the sequel.

We first define cones $\sigma_{S_{\bullet}}$ in $N_{\mathbb{R}}$ that refine the normal fan to the matroid polytope. That means that every cone $\sigma_{S_{\bullet}}$ will be contained in a cone $\sigma_{P(M)_w}$ of the normal fan of the matroid polytope. These cones will therefore have the property that if $w_1, w_2 \in \sigma_{S_{\bullet}}^{\circ}$, then $M_{w_1} = M_{w_2}$.

For a subset $S \subset E$, let e_S be the vector

$$e_S = \sum_{i \in S} e_i$$

in $N_{\mathbb{R}}$. Note that $e_E = 0$. For a flag of subsets,

$$S_{\bullet} = \{\emptyset \subsetneq S_1 \subsetneq \dots \subsetneq S_k \subsetneq E\},$$

let $\sigma_{S_{\bullet}} = \text{Span}_{\geq 0}(e_{S_1}, \dots, e_{S_l})$. If w is in the relative interior of $\sigma_{S_{\bullet}}$ then $w_i < w_j$ if and only if i occurs in an S_k with a larger index than j does. Similarly, given $w \in \mathbb{R}_{\geq 0}^E$, we may pick subsets $\emptyset \subsetneq S_1 \subsetneq \dots \subsetneq S_l \subsetneq E$ such that the elements of E with maximum w -weight are exactly the elements of S_1 , the ones with the next smallest w -weight are the elements of S_2 , and so on. Consequently, w is an element of the relative interior of $\sigma_{S_{\bullet}}$.

Lemma 11.3. *Let M be a matroid. For any $S_\bullet, \sigma_{S_\bullet}$ is contained in a cone of the normal fan to the matroid polytope, $N(P(M))$.*

Proof. It suffices to show that the relative interior $\sigma_{S_\bullet}^\circ$ is contained in the relative interior of a cone of $N(P(M))$. By our description of the normal fan of the matroid polytope, we need to show that for $w_1, w_2 \in \sigma_{S_\bullet}^\circ, M_{w_1} = M_{w_2}$. However, the condition $w_1, w_2 \in \sigma_{S_\bullet}^\circ$ implies the relative ordering of weights of bases with respect to w_1 and w_2 are the same. \square

Lemma 11.4 ([3]). *A cone σ_{S_\bullet} consists of valid vectors w if and only if S_\bullet is a flag of flats.*

Proof. Because the underlying set of the Bergman fan is closed, it suffices to prove the same statement with σ_{S_\bullet} replaced by $\sigma_{S_\bullet}^\circ$.

Suppose some S_i is not a flat of M . Then there exists $j \in \text{cl } S_i \setminus S_i$ and an $i' > i$ such that $j \in S_{i'} \setminus S_{i'-1}$. Because j is in the closure of a flat in $S_{i'-1}$, j must be a loop in $M|_{S_{i'}/S_{i'-1}}$, hence in M_w by the behavior of loops under direct sum.

Now, suppose that each S_i is a flat of M . Let $j \in E$. We must find a basis of M_w containing j . There is an i such that $j \in S_i \setminus S_{i-1}$. Because M has no loops, j is not a loop of $M|_{S_i}$. Because j is not in the closure of S_{i-1} , it is not a loop of $M|_{S_i/S_{i-1}}$, and therefore, it is not a loop in M_w . \square

Lemma 11.4 shows us that the underlying set of the Bergman fan is a union of cones of the form σ_{S_\bullet} . We use this fact to define the *Bergman fan* Δ_M as the simplicial fan in $N_{\mathbb{R}}$ given by σ_{F_\bullet} for flags of flats. A *k-step flag of proper flats* is a sequence of proper flats ordered by containment:

$$F_\bullet = \{\emptyset \subsetneq F_1 \subsetneq \dots \subsetneq F_k \subsetneq E\}$$

The cone associated with F_\bullet is the nonnegative span

$$\sigma_{F_\bullet} = \text{Span}_{\geq 0}\{e_{F_1}, \dots, e_{F_k}\}.$$

Definition 11.5. The *Bergman fan* of M is the fan Δ_M consisting of the cones σ_{F_\bullet} for all flags of flats F_\bullet .

Note that this is a fan structure, not just an underlying set. Ardila and Klivans introduced this fan in [3] and called it the fine subdivision of the Bergman fan of the matroid. Because every flag of flats in a matroid can be extended to a maximal flag of proper flats of length d , the fan Δ_M is pure of dimension d .

Recall that the order complex of a finite poset [91] is the simplicial complex whose vertices are elements of the poset and whose simplices are the chains of the poset. The Bergman fan is a realization of the cone over the order complex of the lattice of proper, non-trivial flats, $L(M) \setminus \{\emptyset, E\}$.

11.2 The Uniform Matroid and the Permutohedral Variety

Of particular interest is the Bergman fan associated with the uniform matroid. We will consider $\mathbb{U} = U_{n+1, n+1}$, the rank $n + 1$ uniform matroid on $E = \{0, \dots, n\}$. Its Bergman fan consists of cones in $N_{\mathbb{R}} = \mathbb{R}^{n+1}/\mathbb{R}$. Every subset of E is a flat. Consequently, the top-dimensional cones of $\Delta_{\mathbb{U}}$ are of the form

$$\text{Span}_{\geq} \{e_{i_0}, e_{i_0} + e_{i_1}, \dots, e_{i_0} + \dots + e_{i_{n-1}}\}$$

for every permutation i_0, \dots, i_n of $0, \dots, n$. Because these cones are generated by a basis of N , $X(\Delta_{\mathbb{U}})$ is smooth. The fan is the barycentric subdivision of the fan corresponding to \mathbb{P}^n , and $X(\Delta_{\mathbb{U}})$ is the toric variety obtained from \mathbb{P}^n by a sequence of blow-ups,

$$\pi_1 : X(\Delta_{\mathbb{U}}) = X_{n-1} \rightarrow \dots \rightarrow X_1 \rightarrow X_0 = \mathbb{P}^n,$$

where $X_{i+1} \rightarrow X_i$ is the blow-up along the proper transforms of the i -dimensional torus-invariant subvarieties of \mathbb{P}^n . The fan, $\Delta_{\mathbb{U}}$ is the normal fan to the permutohedron in its affine span, where the permutohedron is the convex hull of all coordinate permutations of the point $(0, 1, \dots, n) \in \mathbb{R}^{n+1}$ [43].

Example 11.6. The Bergman fan of $U_{d+1, n+1}$ also has a simple description. The top-dimensional cones are of the form

$$\text{Span}_{\geq 0} \{e_{i_0}, e_{i_0} + e_{i_1}, \dots, e_{i_0} + \dots + e_{i_{d-1}}\}$$

for every d -tuple (i_0, \dots, i_{d-1}) of distinct elements of $\{0, \dots, n\}$.

In particular, if $d + 1 = 1$, then the Bergman fan consists simply of the origin. If $d + 1 = 2, n + 1 = 3$, then the Bergman has top-dimensional cones $\mathbb{R}_{\geq 0}e_0, \mathbb{R}_{\geq 0}e_1, \mathbb{R}_{\geq 0}e_2$, and its underlying set is the much-pictured tropical line in the plane with vertex at the origin.

Definition 11.7. The n -dimensional *permutohedral variety* is $X(\Delta_{\mathbb{U}})$, the toric variety associated with the Bergman fan of the uniform matroid $\mathbb{U} = U_{n+1, n+1}$.

Now, the permutohedral variety $X(\Delta_{\mathbb{U}})$ possesses two maps, π_1, π_2 , to \mathbb{P}^n that will be of particular importance. As noted above, every cone of $\Delta_{\mathbb{U}}$ is contained in a cone of $\Delta_{\mathbb{P}^n}$. The induced map $\pi_1 : X(\Delta_{\mathbb{U}}) \rightarrow \mathbb{P}^n$ is the blow-up described above. Now, let $-\Delta_{\mathbb{P}^n}$ be the fan whose cones are of the form $-\sigma$ for each $\sigma \in \Delta_{\mathbb{P}^n}$. The toric variety $X(-\Delta_{\mathbb{P}^n})$ is \mathbb{P}^n but with its torus action precomposed by taking inverse. Now, $\Delta_{\mathbb{U}}$ is a refinement of $-\Delta_{\mathbb{P}^n}$ according to the following lemma:

Lemma 11.8. *Each cone of $\Delta_{\mathbb{U}}$ is contained in a cone of $-\Delta_{\mathbb{P}^n}$.*

Proof. For a subset $S \subseteq E = \{0, 1, \dots, n\}$, $e_S = -e_{E \setminus S}$. Consequently, if

$$F_{\bullet} = \{\emptyset \subsetneq F_1 \subsetneq \dots \subsetneq F_k \subsetneq E\}$$

is a flag of flats, so is

$$F_{\bullet}^c = \{\emptyset \subsetneq E \setminus F_k \subsetneq \dots \subsetneq E \setminus F_1 \subsetneq E\}.$$

Therefore, $-\sigma_{F_{\bullet}} = \sigma_{F_{\bullet}^c}$ is a cone of $\Delta_{\mathbb{U}}$, and $-\Delta_{\mathbb{U}} = \Delta_{\mathbb{U}}$. Since $\Delta_{\mathbb{U}}$ refines $\Delta_{\mathbb{P}^n}$, the conclusion follows. \square

We set π_2 to be the induced birational morphism $\pi_2 : X(\Delta_{\mathbb{U}}) \rightarrow X(-\Delta_{\mathbb{P}^n})$. The existence of π_1 and π_2 shows that the permutohedral variety resolves the indeterminacy of the standard generalized Cremona transformation

$$\text{Crem} : \mathbb{P}^n \dashrightarrow \mathbb{P}^n, \quad [z_0 : \dots : z_n] \mapsto [z_0^{-1} : \dots : z_n^{-1}].$$

This map is induced by multiplication by -1 on N . This map is only a rational map on \mathbb{P}^n because multiplication by -1 does not take cones of $\Delta_{\mathbb{P}^n}$ to cones of $\Delta_{\mathbb{P}^n}$. However, it does induce a birational automorphism of $X(\Delta_{\mathbb{U}})$ fitting into a commutative diagram

$$\begin{array}{ccc} X(\Delta_{\mathbb{U}}) & \xrightarrow{\text{Crem}} & X(\Delta_{\mathbb{U}}) \\ \downarrow \pi_1 & & \downarrow \pi_1 \\ \mathbb{P}^n & \overset{\text{Crem}}{\dashrightarrow} & \mathbb{P}^n. \end{array}$$

Note that $\pi_2 = \pi_1 \circ \text{Crem}$. There is a natural morphism $\pi_1 \times \pi_2 : X(\Delta_{\mathbb{U}}) \rightarrow \mathbb{P}^n \times \mathbb{P}^n$. The presence of the Cremona transformation is related to the use of reciprocal hyperplanes as in [78, 84].

There are natural subvarieties of $X(\Delta_{\mathbb{U}})$, called proper transforms, that are associated with subspaces of \mathbb{P}^n :

Definition 11.9. Let V be a $(d + 1)$ -dimensional subspace of \mathbf{k}^{n+1} that is not contained in any hyperplane, and let $\mathbb{P}(V) \subset \mathbb{P}^n$ be its projectivization. The *proper transform* of $\mathbb{P}(V)$ in $X(\Delta_{\mathbb{U}})$ is

$$\widetilde{\mathbb{P}(V)} = \overline{\pi_1^{-1}(\mathbb{P}(V)^\circ)}.$$

By the blow-up interpretation of the permutohedral variety, the proper transform is an iterated blow-up of $\mathbb{P}(V)$ at its intersections with the coordinate subspaces of \mathbb{P}^n . First one blows up the 0-dimensional intersections, then the proper transforms

of 1-dimensional intersections, and then so on. It is easily seen that $\widetilde{\mathbb{P}(V)}$ intersects the torus orbits of $X(\Delta_{\mathbb{U}})$ properly.

There are two piecewise linear functions of $\Delta_{\mathbb{U}}$ that will be of great importance in the sequel. We have the piecewise linear function on $\Delta_{\mathbb{P}^n}$ from Example 10.12,

$$\alpha = \min(0, w_1, \dots, w_n)$$

We will abuse notation and denote the pullback $\pi_1^* \alpha$ on $X(\Delta_{\mathbb{U}})$ by α . We can also pull back the analogous piecewise linear function from $X(-\Delta_{\mathbb{P}^n})$ to obtain

$$\beta = \min(0, -w_1, \dots, -w_n).$$

The non-equivariant cohomology class $\iota^* \alpha$ corresponds to the proper transform of a generic hyperplane in $X(\Delta_{\mathbb{U}})$. On the other hand, $\iota^* \beta$ corresponds to the closure in $X(\Delta_{\mathbb{U}})$ of reciprocal hyperplanes in $(\mathbf{k}^*)^n$ given by

$$a_0 + \frac{a_1}{x_1} + \dots + \frac{a_n}{x_n} = 0$$

for generic choices of $a_i \in \mathbf{k}$ where $(\mathbf{k}^*)^n$ is the big open torus. We may suppress ι^* and write α and β for the induced non-equivariant cohomology classes.

11.3 The Bergman Fan as a Minkowski Weight

In this section, we show how the Bergman fan Δ_M can be thought of a Minkowski weight on $\Delta_{\mathbb{U}}$. We will let M be a rank $d + 1$ matroid so that for F_{\bullet} , a flag of flats in M of maximum length, $\sigma_{F_{\bullet}}$ is a d -dimensional cone.

Definition 11.10. The Minkowski weight on $\Delta_{\mathbb{U}}$ corresponding to the rank $d + 1$ matroid M is given by $\Delta_M : \Delta_{\mathbb{U}}^{(n-d)} \rightarrow \mathbb{Z}$ where

$$\Delta_M(\sigma_{F_{\bullet}}) = \begin{cases} 1 & \text{if } F_{\bullet} \text{ is a flag of flats in } M \\ 0 & \text{otherwise.} \end{cases}$$

Lemma 11.11. *The function Δ_M is a Minkowski weight on $\Delta_{\mathbb{U}}$*

Proof. Let $\sigma_{G_{\bullet}}$ be a codimension 1 cone in Δ_M . Then $\sigma_{G_{\bullet}}$ corresponds to a flag of flats of length $d - 1$ of the form

$$G_{\bullet} = \{\emptyset \subsetneq F_1 \subsetneq \dots \subsetneq F_{j-1} \subsetneq F_{j+1} \subsetneq \dots \subsetneq F_d\}$$

where $r(F_i) = i$. We restrict M to F_{j+1} since any cone containing $\sigma_{G_{\bullet}}$ corresponds to a flag of flats refining G_{\bullet} . The flats F'_1, \dots, F'_i of $M|_{F_{j+1}}$ properly containing F_{j-1}

are exactly the flats that can be inserted in G_\bullet to obtain a flag of flats of length d . The sets $F'_k \setminus F_{j-1}$ partition $F_{j+1} \setminus F_{j-1}$ by Definition 3.5. Let $(F_\bullet)_k$ be the flag of flats given by inserting F'_k into G_\bullet . Therefore, $u_{\sigma_{(F_\bullet)_k}/\sigma_{G_\bullet}} = e_{F'_k}$. From

$$\sum_{k=1}^l \Delta_M(\sigma_{(F_\bullet)_k}) u_{\sigma_{(F_\bullet)_k}/\sigma_{G_\bullet}} = \sum_{k=1}^l e_{F'_k} \in \text{Span}(e_{F_{j-1}}, e_{F_{j+1}}) \subseteq N_{\sigma_{G_\bullet}},$$

we see that Δ_M is a Minkowski weight. □

The following is straightforward:

Lemma 11.12. *Let $V \subset \mathbf{k}^{n+1}$ be a representation of a simple rank $d + 1$ matroid on $\{0, 1, \dots, n\}$. Then the Minkowski weight Δ_M is the associated cocycle of the proper transform $\widetilde{\mathbb{P}(V)} \subset X(\Delta_{\mathbb{U}})$.*

It follows that the Bergman fan is the same thing as the tropicalization of $\mathbb{P}(V)^\circ$. In less technical terms, the Minkowski weight records the toric strata that $\widetilde{\mathbb{P}(V)}$ intersects. These corresponds to flags of flats of M .

It turns out that among tropicalizations, the Bergman fans exactly correspond to the tropicalizations of linear subspaces. This follows from a theorem that emerged in a discussion of Mikhalkin and Ziegler and which was written down in [53]. See [45] for another proof perhaps more suitable for the statement here. We call this theorem, “the duck theorem” in the sense that *if it looks like a duck, swims like a duck, and quacks like a duck, then it is probably a duck*:

Theorem 11.13. *Let $Y^\circ \subset (\mathbf{k}^*)^n \subset X(\Delta_{\mathbb{U}})$ be a variety whose associated cocycle is Δ_M , the Bergman fan of a matroid M . Then $Y^\circ = \mathbb{P}(V)^\circ$ where V is a subspace realizing the matroid M .*

This theorem can be used to come up with counterexamples to questions about tropical lifting, that is, to determine whether a Minkowski weight is the associated cocycle of an algebraic subvariety of $(\mathbf{k}^*)^n$. By starting with a non-representable matroid, one can produce a Minkowski weight Δ_M on $X(\Delta_{\mathbb{U}})$ which is not the associated cocycle of any subvariety of $X(\Delta_{\mathbb{U}})$. Indeed if such a subvariety existed, it would be the closure of $\mathbb{P}(V)^\circ$ where V is a representation of M .

In addition, this theorem can be used to study which homology classes in $X(\Delta_{\mathbb{U}})$ can be realized by irreducible subvarieties. In [45], it is shown that the Bergman class of a loopless matroid of rank $d + 1$, considered as cohomology class is nef in the sense that it has nonnegative intersection with all effective cycles of complementary dimension. Moreover, it generates an extremal ray of the nef cone in dimension d . It is effective in that it can be written as the sum of effective cycles. However, it can be represented by an irreducible subvariety over \mathbf{k} if and only if the matroid is representable over \mathbf{k} . This shows that the study of the realizability of homology classes in toric varieties is heavily dependent on the ground field and is as complicated as the theory of representability of matroids.

11.4 (r_1, r_2) -Truncation of Bergman fans

One advantage of working with Minkowski weights in Δ_{\cup} rather than matroids is that they allow a new operation, (r_1, r_2) -truncation as developed by Huh and Katz. More details can be found in [45].

Definition 11.14. Let M be a matroid of rank $d + 1$. Let r_1, r_2 be integers with $1 \leq r_1 \leq r_2 \leq d + 1$. The (r_1, r_2) -truncation of M is the Minkowski weight of dimension $r_2 - r_1 + 1$, $\Delta_{M[r_1, r_2]}$ on Δ_{\cup} defined as follows: if $\sigma_{F_{\bullet}}$ is the cone determined by

$$F_{\bullet} = \{F_{r_1} \subsetneq F_{r_1+1} \subsetneq \dots \subsetneq F_{r_2}\}$$

then $\Delta_{M[r_1, r_2]}(\sigma_{F_{\bullet}})$ is

- (1) $|\mu(\emptyset, F_{r_1})|$ if each F_i is a flat of M of rank i , or
- (2) 0 otherwise.

Note that the $(1, r_2)$ -truncation of M is the Bergman fan of the matroid $\text{Trunc}^{r_2+1}(M)$. However, if $r_1 \geq 2$, the (r_1, r_2) -truncation of M is not the Bergman fan of any matroid. It is not obvious that $\Delta_{M[r_1, r_2]}$ is indeed a Minkowski weight. This will follow from the following lemma whose proof we will defer until Sect. 12.2.

Proposition 11.15. *We have the following relation among Minkowski weights:*

$$\alpha^{d-r_2} \beta^{r_1-1} \cup \Delta_M = \Delta_{M[r_1, r_2]}.$$

Moreover, we have the following degree computations:

- (1) $\deg(\alpha \cup \Delta_{M[r, r]}) = \mu^{r-1}$
- (2) $\deg(\beta \cup \Delta_{M[r, r]}) = \mu^r$

where μ^r is a coefficient of the reduced characteristic polynomial.

Note that applying powers of α in the above makes geometric sense. The theorem states that $\alpha^k \cup \Delta_M = \Delta_{\text{Trunc}^{d+1-k}(M)}$. Because α is the class of a hyperplane in \mathbb{P}^n , we are intersecting a projective subspace $\mathbb{P}(V)$ with generic hyperplanes to obtain a projective subspace whose matroid is a truncation of M by Lemma 5.7.

We have the following easy corollary:

Corollary 11.16. *There is the Following intersection-theoretic description of coefficients of the reduced characteristic polynomial:*

$$\deg(\alpha^{d-r} \beta^r \cup \Delta_M) = \mu^r.$$

Proof. By applying the above proposition, we have

$$\deg(\alpha^{d-r} \beta^r \cup \Delta_M) = \deg(\beta \cup (\alpha^{d-r} \beta^{r-1} \cup \Delta_M)) = \deg(\beta \cup \Delta_{M[r, r]}) = \mu^r.$$

□

11.5 The Bergman Fan as a Cryptomorphic Definition of a Matroid

Minkowski weights on $\Delta_{\mathbb{U}}$ are interesting combinatorial objects in their own right. They contain matroids on $\{0, 1, \dots, n\}$ as a specific subclass but are closed under the truncation operation defined above. It would be an interesting exercise to generalize the matroid operations that we discussed in Sect. 5 to them. Here, extensions are closely related to the notion of tropical modifications as introduced by Mikhalkin. See [87] for more on tropical modifications.

It is a consequence of Proposition 11.15 that if c is a Minkowski weight associated with a rank $d + 1$ matroid on $\{0, 1, \dots, n\}$ then $\alpha^d \cup c$ is the Minkowski weight corresponding to the rank 0 matroid. Consequently, $\deg(\alpha^d \cup c) = 1$. The converse is true by the following theorem which was noted by Mikhalkin, Sturmfels, and Ziegler as described by [71] and proved by Fink [34].

Theorem 11.17. *Let $c \in A^{n-d}(X(\Delta_{\mathbb{U}}))$ be a Minkowski weight of codimension $n - d$. Then $c = \Delta_M$ for some matroid M of rank $d + 1$ if and only if $\deg(\alpha^d \cup c) = 1$.*

The above theorem can be thought of as cryptomorphic definition of a matroid. In the course of the proof, Fink gives an explicit recipe for finding the matroid polytope associated with c .

12 Log-Concavity of the Characteristic Polynomial

12.1 Proof of Log-Concavity

We will prove that $\chi_M(q)$ is log-concave when M is a representable matroid. By an easy algebra computation, it suffices to show that the reduced characteristic polynomial $\bar{\chi}_M(q)$ is log-concave.

We will identify μ^k with intersection numbers on a particular algebraic variety. Let $V \subset \mathbf{k}^{n+1}$ be a linear subspace representing M . We will give the proof assuming Proposition 11.15 for now.

Proposition 12.1. *There is a complete irreducible algebraic variety $\widetilde{\mathbb{P}(V)}$ with nef divisors α, β such that*

$$\mu^k = \deg(\alpha^{d-k} \beta^k \cap [\widetilde{\mathbb{P}(V)}]).$$

Log-concavity then follows by applying the Khovanskii–Teissier inequality [61, Example 1.6.4]:

Theorem 12.2. *Let X be a complete irreducible d -dimensional variety, and let α, β be nef divisors on X . Then $\deg((\alpha^{d-k} \beta^k) \cap [X])$ is a log-concave sequence.*

Proof. We give an outline of the proof of this inequality. By resolution of singularities, Chow’s lemma and the projection formula, we may suppose that X is smooth and projective. It suffices to show

$$\deg((\alpha^{d-k}\beta^k) \cap [X])^2 \geq \deg((\alpha^{d-k+1}\beta^{k-1}) \cap [X]) \deg((\alpha^{d-k-1}\beta^{k+1}) \cap [X]).$$

By Kleiman’s criterion [56] and continuity of intersection numbers, we can perturb α and β in $A^1(X)$ so that they are ample classes. By homogeneity of the desired inequality, we can replace α and β by $m\alpha$ and $m\beta$ for $m \in \mathbb{Z}_{\geq 1}$ so that they are very ample. By the Kleiman–Bertini theorem [57], the intersection $\alpha^{d-k-1}\beta^{k-1} \cap [X]$ is represented by an irreducible smooth surface Y . Now, we need only prove

$$\deg(\alpha\beta \cap [Y])^2 \geq \deg(\alpha^2 \cap [Y]) \deg(\beta^2 \cap [Y]).$$

By the Hodge index theorem, the intersection product on $A^1(Y)$ has signature $(1, n - 1)$. The restriction of the intersection product to the subspace spanned by α and β is indefinite. The desired inequality is exactly the non-positivity of the determinant of the matrix of the intersection product. \square

We now give the proof of Proposition 12.1:

Proof. Let M be represented by the linear subspace $V \subset \mathbf{k}^{n+1}$, and let $\widetilde{\mathbb{P}(V)}$ be the proper transform of $\mathbb{P}(V)$ in $X(\Delta_{\mathbb{U}})$. Let α, β be as in Sect. 11.1. Now, we know

$$\Delta_M \cap [X(\Delta_{\mathbb{U}})] = \widetilde{\mathbb{P}(V)}$$

by Lemmas 10.14 and 11.12. Consequently,

$$\deg(\alpha^{d-k}\beta^k \cap [\widetilde{\mathbb{P}(V)}]) = \deg(\alpha^{d-k}\beta^k \cup \Delta_M) = \mu^k$$

where the last equality follows from Corollary 11.16. \square

By viewing α and β as hyperplane classes on \mathbb{P}^n , we immediately have the following corollary:

Corollary 12.3. *The cycle class of $[(\pi_1 \times \pi_2)(\widetilde{\mathbb{P}(V)})]$ in $\mathbb{P}^n \times \mathbb{P}^n$ is*

$$\mu^0[\mathbb{P}^d \times \mathbb{P}^0] + \mu^1[\mathbb{P}^{d-1} \times \mathbb{P}^1] + \cdots + \mu^r[\mathbb{P}^0 \times \mathbb{P}^d].$$

Now, we outline Lenz’s proof of Mason’s conjecture in the representable case which relates the f -vector of a matroid to the characteristic polynomial of a different, well-chosen matroid. We form the f -polynomial of M by setting

$$f_M(q) = \sum_{i=0}^{d+1} f_i q^{d+1-i}.$$

The free co-extension of M is the matroid

$$M \times e = (M^* +_E e)^*$$

for a new element e . Because M is representable, after a possible extension of the field \mathbf{k} , so is $M \times e$. Lenz proves the following formula for the characteristic polynomial of $M \times e$

$$(-1)^d \chi_{M \times e}(-q) = (1 + q)f_M(q)$$

by considering various specializations of the rank-generating polynomial. The log-concavity of f_M then follows from the log-concavity of the reduced characteristic polynomial of $M \times e$.

12.2 Intersection Theory Computations

Now, we will give a proof of Proposition 11.15 following the exposition of [45]. We will derive the conclusion from three lemmas:

Lemma 12.4. *Let M be a rank $d + 1$ matroid on $E = \{0, \dots, n\}$. Then,*

$$\alpha \cup \Delta_M = \Delta_{M[1,d-1]}$$

in $A^*(X(\Delta_{\mathbb{U}}))$.

Note here that $\Delta_{M[1,d-1]} = \Delta_{\text{Trunc}^d(M)}$.

Lemma 12.5. *Let M be a rank $d + 1$ matroid on $E = \{0, \dots, n\}$. Then,*

$$\beta \cup \Delta_{M[r_1,d]} = \Delta_{M[r_1+1,d]}$$

in $A^*(X(\Delta_{\mathbb{U}}))$.

To simplify the proofs of these lemmas, we homogenize our piecewise linear functions. We begin with the lattice \mathbb{Z}^{n+1} spanned by e_0, e_1, \dots, e_n . Let w_0, w_1, \dots, w_n be the induced coordinates on $\mathbb{R}^{n+1} = \mathbb{Z}^{n+1} \otimes \mathbb{R}$. We will let our lattice N be defined by $w_0 = 0$ in \mathbb{Z}^{n+1} . Consider the quotient by the diagonal line

$$p : \mathbb{R}^{n+1} \rightarrow \mathbb{R}^{n+1}/\mathbb{R} \cong N_{\mathbb{R}}.$$

Let Δ' be the fan in \mathbb{R}^{n+1} whose cones are the inverse images of the cones of $\Delta_{\mathbb{U}}$ under p . Write $\sigma'_{F_{\bullet}} = p^{-1}(\sigma_{F_{\bullet}})$. By assigning it the same values on relevant cones, we may treat Δ_M as a Minkowski weight on Δ' , denoted by Δ'_M . We will consider the piecewise linear functions on Δ' ,

$$\begin{aligned} \alpha' &= \min(w_0, w_1, \dots, w_n) \\ \beta' &= \min(-w_0, -w_1, \dots, -w_n) \end{aligned}$$

Then α and β are the restrictions of α' and β' to $N_{\mathbb{R}}$, respectively. Because the entire situation is invariant under translation by the diagonal vector $e_0 + e_1 + \dots + e_n$, we can do the intersection theory computation on \mathbb{R}^{n+1} and then intersect with $N_{\mathbb{R}}$.

We first give the proof of Lemma 12.4.

Proof. The Minkowski weight $\alpha \cup \Delta'_M$ is supported on d -dimensional cones in $\Delta'_{\mathbb{U}}$. They correspond to $(d - 1)$ -step flags of proper flats

$$F_{\bullet} = \{\emptyset \subsetneq F_1 \subsetneq F_2 \subsetneq \dots \subsetneq F_{d-1} \subsetneq F_d = E\}.$$

The cone $\sigma'_{F_{\bullet}}$ is contained in $\sigma'_{G_{\bullet}}$ if and only if the flag G_{\bullet} is obtained from F_{\bullet} by inserting a single flat. Write this relation as $G_{\bullet} > F_{\bullet}$. This flat must be inserted between two flats $F_j \subset F_{j+1}$ where $r(F_{j+1}) = r(F_j) + 2$. There is a unique choice of j where this happens. Suppose G_{\bullet} is obtained from inserting a proper flat F between $F_j \subset F_{j+1}$. Let $u_{G_{\bullet}/F_{\bullet}}$ be an integer vector in $\sigma_{G_{\bullet}}$ that generates the image of $\sigma_{G_{\bullet}}$ in $N/N_{\sigma_{F_{\bullet}}}$. We may choose $u_{G_{\bullet}/F_{\bullet}}$ to be e_F . The value of $\alpha' \cup \Delta'_M$ on $\sigma'_{F_{\bullet}}$ is given by

$$(\alpha' \cup \Delta'_M)(\sigma'_{F_{\bullet}}) = - \sum_{G_{\bullet} > F_{\bullet}} \alpha'_{G_{\bullet}}(u_{G_{\bullet}/F_{\bullet}}) + \alpha'_{F_{\bullet}}\left(\sum_{G_{\bullet} > F_{\bullet}} u_{G_{\bullet}/F_{\bullet}}\right)$$

where $\alpha'_{G_{\bullet}}$ (respectively $\alpha'_{F_{\bullet}}$) is the linear function on $N_{\sigma_{G_{\bullet}}}$ (on $N_{\sigma_{F_{\bullet}}}$) which equals α' on $\sigma_{G_{\bullet}}$ (on $\sigma_{F_{\bullet}}$).

We now compute the right side. Because $F \neq E$,

$$\alpha'_{G_{\bullet}}(u_{G_{\bullet}/F_{\bullet}}) = \alpha'(e_F) = 0.$$

Let f be the number of flats that can be inserted between F_j and F_{j+1} . Because every element of $F_{j+1} \setminus F_j$ is contained in exactly one such flat F by Definition 3.5 (3.5), we have

$$\sum_{G_{\bullet} > F_{\bullet}} u_{G_{\bullet}/F_{\bullet}} = e_{F_{j+1}} + (f - 1)e_{F_j}.$$

and so

$$\alpha'_{F_{\bullet}}\left(\sum_{G_{\bullet} > F_{\bullet}} u_{G_{\bullet}/F_{\bullet}}\right) = \begin{cases} 1 & \text{if } F_{j+1} = E \\ 0 & \text{otherwise} \end{cases}$$

Consequently, we have

$$(\alpha' \cup \Delta'_M)(\sigma'_{F_\bullet}) = \begin{cases} 1 & \text{if } j = d - 1 \\ 0 & \text{if otherwise.} \end{cases}$$

Therefore $\alpha \cup \Delta_M$ is non-zero on exactly the cones in the support of $\Delta_{\text{Trunc}^d(M)}$ where it takes the value 1. \square

Now, we prove Lemma 12.5.

Proof. The setup is as in the proof of the above lemma. For a flag of flats

$$F_\bullet = \{\emptyset = F_0 \subsetneq F_1 \subsetneq F_2 \subsetneq \cdots \subsetneq F_{d-1} \subsetneq F_d = E\},$$

write $v_{F_\bullet} = |\mu(\emptyset, F_1)|$. Let F be a flat inserted between $F_j \subset F_{j+1}$ to obtain a flag of flats G_\bullet . Here, we have

$$(\beta' \cup \Delta'_{M[r_1, d]})(\sigma'_{F_\bullet}) = - \sum_{G_\bullet \triangleright F_\bullet} v_{G_\bullet} \beta'_{G_\bullet}(u_{G_\bullet/F_\bullet}) + \beta'_{F_\bullet} \left(\sum_{G_\bullet \triangleright F_\bullet} v_{G_\bullet} u_{G_\bullet/F_\bullet} \right)$$

Because $F \neq \emptyset$,

$$\beta'_{G_\bullet}(u_{G_\bullet/F_\bullet}) = \beta'(e_F) = -1.$$

Now we consider two cases: $F_j \neq \emptyset$ and $F_j = \emptyset$. If $F_j \neq \emptyset$,

$$\sum_{G_\bullet \triangleright F_\bullet} v_{G_\bullet} \beta'_{G_\bullet}(u_{G_\bullet/F_\bullet}) = f |\mu(\emptyset, F_1)|.$$

and

$$\sum_{G_\bullet \triangleright F_\bullet} v_{G_\bullet} u_{G_\bullet/F_\bullet} = |\mu(\emptyset, F_1)| (e_{F_{j+1}} + (f - 1)e_{F_j}).$$

Therefore, $\beta'(\sum_{G_\bullet \triangleright F_\bullet} v_{G_\bullet} u_{G_\bullet/F_\bullet}) = f |\mu(\emptyset, F_1)|$ and $(\beta' \cup \Delta'_{M[r_1, d]})(\sigma'_{F_\bullet}) = 0$.

If $F_j = \emptyset$,

$$\sum_{G_\bullet \triangleright F_\bullet} v_{G_\bullet} \beta'(u_{G_\bullet/F_\bullet}) = - \sum_{F \triangleleft F_1} |\mu(\emptyset, F)|$$

and

$$\beta'_{F_\bullet} \left(\sum_{G_\bullet \triangleright F_\bullet} v_{G_\bullet} u_{G_\bullet/F_\bullet} \right) = \beta' \left(\sum_{F \triangleleft F_1} |\mu(\emptyset, F)| e_F \right) = - \sum_{a \in F \triangleleft F_1} |\mu(\emptyset, F)|$$

where some $a \in F_1$ chosen so to maximize the quantity on the right. Then,

$$(\beta' \cup \Delta_{M[r_1,d]})(\sigma'_{F_\bullet}) = \sum_{a \notin F \triangleleft F_1} |\mu(\emptyset, F)| = |\mu(\emptyset, F_1)|$$

where the last equality follows from Lemma 7.11. Therefore $\beta \cup \Delta_M$ is non-zero on exactly the cones in the support of $\Delta_{M[r_1+1,d]}$ where it takes the expected value. \square

Lemma 12.6. *We have the following equality of degrees:*

- (1) $\deg(\alpha \cup \Delta_{M[r,r]}) = \mu^{r-1}$
- (2) $\deg(\beta \cup \Delta_{M[r,r]}) = \mu^r$

where μ^k is the coefficient of the reduced characteristic polynomial.

Proof. Let γ' be α' or β' as above. Because the only codimension 1 cone in $\Delta_{M[r,r]}$ is the origin, we have the following formula for the degree:

$$\deg(\gamma' \cup \Delta_{M[r,r]}) = - \sum_{F \in L(M)_r} |\mu(\emptyset, F)| \gamma'(e_F) + \gamma' \left(\sum_{F \in L(M)_r} |\mu(\emptyset, F)| e_F \right)$$

For $\gamma' = \alpha'$, this becomes

$$\alpha' \left(\sum_{F \in L(M)_r} |\mu(\emptyset, F)| e_F \right) = \sum_{a \in F \in L(M)_r} |\mu(\emptyset, F)| = \mu^{r-1}$$

for some $a \in E$ where the last equality follows from Lemma 7.15. For $\gamma' = \beta'$, we have

$$\deg(\beta' \cup \Delta_{M[r,r]}) = \sum_{F \in L(M)_r} |\mu(\emptyset, F)| - \sum_{a \in F \in L(M)_r} |\mu(\emptyset, F)| = \sum_{a \notin F \in L(M)_r} |\mu(\emptyset, F)| = \mu^r$$

\square

13 Future Directions

One would like to prove the log-concavity of the characteristic polynomial in the non-representable case. There are a couple of lines of attack that are being considered in future work by Huh individually and with the author.

The proof presented above requires representability to invoke the Khovanskii–Teissier inequality. The proof of the Khovanskii–Teissier inequality reduces to the Hodge index theorem on a particular algebraic surface. Recall that the log-concavity statement concerns only three consecutive coefficients of the characteristic polynomial at a time. These three coefficients are intersection numbers on this surface.

One might try to prove a combinatorial analogue of the Hodge index theorem on the combinatorial analogue of this surface which is the truncated Bergman fan $\Delta_{M[r,r+1]}$ by showing that a *combinatorial intersection matrix* has a single positive eigenvalue. This combinatorial intersection matrix is called the Tropical Laplacian and will be investigated in future work with Huh. Unfortunately, the algebraic geometric arguments used in proofs of the Hodge index theorem do not translate into combinatorics, and the hypotheses for a combinatorial Hodge index theorem are unclear. Still, the conclusion of the combinatorial Hodge index theorem for $\Delta_{M[r,r+1]}$ has been verified experimentally for all matroids on up to nine elements by Theo Belaire [7] using the matroid database of Mayhew and Royle [66].

Another approach to the Rota–Heron–Welsh conjecture is to relax the definition of representability. In our proof above, we needed the Chow cohomology class Δ_M to be Poincaré-dual to an irreducible subvariety of $X(\Delta_{\mathbb{U}})$ in order to apply the Khovanskii–Teissier inequality. However, it is sufficient that some positive integer multiple of Δ_M be Poincaré-dual to an irreducible subvariety. It is part of a general philosophy of Huh [45] that it is very difficult to understand which homology classes are representable while it is significantly easier to understand their cone of positive multiples.

Another question of interest is to understand the relation between the work of Huh–Katz and Fink–Speyer. What does positivity (in the sense of [61]) say about K -theory. How does positivity restrict the Tutte polynomial? Are there other specializations of the Tutte polynomial that obey log-concavity?

Finally, the author would like to promote the importance of a combinatorial study of Minkowski weights on the permutohedral variety. They are very slight enlargements of the notion of matroids. One can certainly introduce notions of deletion and contraction and therefore minors. Are there interesting structure theorems? To this author, (r_1, r_2) -truncation is an attractive and useful operation and should be situated in a general combinatorial theory.

References

1. Aigner, M.: Whitney numbers. In: Combinatorial Geometries, pp. 139–160. Encyclopedia of Mathematics and Its Applications, vol. 29. Cambridge University Press, Cambridge (1987)
2. Aluffi, P.: Grothendieck classes and Chern classes of hyperplane arrangements. *Int. Math. Res. Not.* **8**, 1873–1900 (2013)
3. Ardila, F., Klivans, C.: The Bergman complex of a matroid and phylogenetic trees. *J. Comb. Theory Ser. B* **96**, 38–49 (2006)
4. Ardila, F., Fink, A., Rincón, F.: Valuations for matroid polytope subdivisions. *Can. J. Math.* **62**, 1228–1245 (2010)
5. Allermann, L., Rau, J.: First steps in tropical intersection theory. *Math. Z.* **264**, 633–670 (2010)
6. Athanasiadis, C.: Characteristic polynomials of subspace arrangements and finite fields. *Adv. Math.* **122**, 193–233 (1996)
7. Belaire, T., Katz, E.: A computer verification of the Hodge index conjecture for small matroids (in preparation)

8. Bergman, G.M.: The logarithmic limit-set of an algebraic variety. *Trans. Am. Math. Soc.* **157**, 459–469 (1971)
9. Bieri, R., Groves, J.R.J.: The geometry of the set of characters induced by valuations. *J. Reine Angew. Math.* **347**, 168–195 (1984)
10. Billera, L., Jia, N., Reiner, V.: A quasisymmetric function for matroids. *Eur. J. Comb.* **30**, 1727–1757 (2009)
11. Birkhoff, G.: A determinant formula for the number of ways of coloring a map. *Ann. Math.* **14**, 42–46 (1912)
12. Bixby, R.E.: On Reid’s characterization of the ternary matroids. *J. Comb. Theory Ser. B.* **26**, 174–204 (1979)
13. Bixby, R.E., Cunningham, W.H.: Matroid optimization and algorithms. *Handbook of Combinatorics*, vol. 551–609. Elsevier, Amsterdam (1995)
14. Björner, A., Las Vergnas, M., Sturmfels, B., White, N., Ziegler, G.: *Oriented Matroids*. *Encyclopedia of Mathematics and Its Applications*, vol. 46. Cambridge University Press, Cambridge (1999)
15. Borovik, A., Gelfand, I.M., White, N.: *Coxeter Matroids*. Birkhäuser, Boston, MA (2003)
16. Brylawski, T.: *Constructions*. *Theory of matroids*, 127–223, *Encyclopedia of Mathematics and Its Applications*, vol. 26. Cambridge University Press, Cambridge (1986)
17. Brylawski, T., Kelly, D.: *Matroids and Combinatorial Geometries*. *Carolina Lecture Series*. Department of Mathematics, University of North Carolina, Chapel Hill, NC (1980)
18. Cartwright, D.: Lifting matroid divisors on tropical curves. *Res. Math. Sci.* **2**, 23 (2015)
19. Cox, D., Little, J., Schenck, H.: *Toric Varieties*. *Graduate Studies in Mathematics*, vol. 124. American Mathematical Society, Providence, RI (2011)
20. Crapo, H.H.: Single-element extensions of matroids. *J. Res. Natl. Bur. Stand. Sect. B* **69**, 55–65 (1965)
21. Crapo, H.H.: The Tutte polynomial. *Aeq. Math.* **3**, 211–229 (1969)
22. Crapo, H.H., Rota, G.-C.: *On the Foundations of Combinatorial Theory: Combinatorial Geometries*, Preliminary edition. MIT, Cambridge, MA, London (1970)
23. Dat, J.-F., Orlik, S., Rapoport, M.: *Period domains over finite and p -adic fields*. *Cambridge Tracts in Mathematics*, vol. 183. Cambridge University Press, Cambridge (2010)
24. Dawson, J.: A collection of sets related to the Tutte polynomial of a matroid. In: *Graph theory, Singapore 1983*, pp. 193–204. *Lecture Notes in Mathematics*, vol. 1073. Springer, Berlin (1984)
25. Denham, G.: Toric and tropical compactifications of hyperplane complements. *Ann. Fac. Sci. Toulouse Math.* **23**, 297–333 (2014)
26. Derksen, H.: Symmetric and quasi-symmetric functions associated to polymatroids. *J. Algebraic Comb.* **30**, 43–86 (2009)
27. Derksen, H., Fink, A.: Valuative invariants for polymatroids. *Adv. Math.* **225**, 1840–1892 (2010)
28. Diestel, R.: *Graph Theory*. Springer, Heidelberg (2010)
29. Dress, A.: Duality theory for finite and infinite matroids with coefficients. *Adv. Math.* **59**, 97–123 (1986)
30. Dress, A., Wenzel, W.: Grassmann-Plücker relations for matroids with coefficients. *Adv. Math.* **86**, 68–110 (1991)
31. Edidin, D., Graham, W.: Equivariant intersection theory. *Invent. Math.* **131**, 595–634 (1998)
32. Edmonds, J.: *Submodular functions, matroids, and certain polyhedra*. *Combinatorial Structures and Their Applications*. *Proceedings of Calgary International Conference*, Calgary, Alberta, 1969, pp. 69–87. Gordon and Breach, New York (1970)
33. Feichtner, E.V., Sturmfels, B.: Matroid polytopes, nested sets, and Bergman fans. *Port. Math.* **62**, 437–468 (2005)
34. Fink, A.: Tropical cycles and Chow polytopes. *Beitr. Algebra Geom.* **54**, 13–40 (2013)
35. Fink, A., Moci, L.: Matroids over a ring. *J. Eur. Math. Soc.* **18**, 681–731 (2016)
36. Fink, A., Speyer, D.: K-classes of matroids and equivariant localization. *Duke Math. J.* **161**, 2699–2723 (2012)

37. Fulton, W.: Intersection Theory. *Ergebnisse der Mathematik und ihrer Grenzgebiete*, vol. 2(3). Springer, Berlin (1984)
38. Fulton, W.: Introduction to Toric Varieties. *Annals of Mathematics Studies*, vol. 131. Princeton University Press, Princeton, NJ (1993)
39. Fulton, W., Sturmfels, B.: Intersection theory on toric varieties. *Topology* **36**(2), 335–353 (1997)
40. Geelen, J., Gerards, B., Whittle, G.: Structure in minor-closed classes of matroids. In: *Surveys in Combinatorics 2013*, pp. 327–362. London Mathematical Society Lecture Notes, vol. 409. Cambridge University Press, Cambridge (2013)
41. Geelen, J., Gerards, B., Whittle, G.: Solving Rota’s conjecture. *Not. Am. Math. Soc.* **61**, 736–743 (2014)
42. Gelfand, I.M., Goresky, M., MacPherson, R., Serganova, V.: Combinatorial geometries, convex polyhedra, and Schubert cells. *Adv. Math.* **63**, 301–316 (1987)
43. Gelfand, I.M., Kapranov, M.M., Zelevinsky, A.V.: *Discriminants, Resultants, and Multidimensional Determinants*. Birkhäuser, Boston (1994)
44. Huh, J.: Milnor numbers of projective hypersurfaces and the chromatic polynomial of graphs. *J. Am. Math. Soc.* **25**, 907–927 (2012)
45. Huh, J.: Rota’s conjecture and positivity of algebraic cycles in permutohedral varieties. Ph.D. thesis, University of Michigan (2014)
46. Huh, J.: h -vectors of matroids and logarithmic concavity. *Adv. Math.* **270**, 49–59 (2015)
47. Huh, J., Katz, E.: Log-concavity of characteristic polynomials and the Bergman fan of matroids. *Math. Ann.* **354**, 1103–1116 (2012)
48. Ingleton, A.W.: Representation of matroids. In: *Combinatorial Mathematics and its Applications*. Proceedings of Conference, Oxford, 1969, pp. 149–167. Academic Press, London (1971)
49. Ingleton, A.W., Main, R.A.: Non-algebraic matroids exist. *Bull. London Math. Soc.* **7**, 144–146 (1975)
50. Jensen, D.: The locus of Brill-Noether general graphs is not dense (2015). Preprint, [arXiv:1405.6338](https://arxiv.org/abs/1405.6338)
51. Katz, E.: A tropical toolkit. *Expo. Math.* **27**, 1–36 (2009)
52. Katz, E.: Tropical intersection theory from toric varieties. *Collect. Math.* **63**, 29–44 (2012)
53. Katz, E., Payne, S.: Realization spaces for tropical fans. In: *Combinatorial Aspects of Commutative Algebra and Algebraic Geometry*, pp. 73–88. Abel Symposium, vol. 6. Springer, Berlin (2011)
54. Katz, E., Stapledon, A.: Tropical geometry and the motivic nearby fiber. *Compos. Math.* **148**, 269–294 (2012)
55. Kinser, R.: New inequalities for subspace arrangements. *J. Combin. Theory Ser. A* **118**, 152–161 (2011)
56. Kleiman, S.L.: Toward a numerical theory of ampleness. *Ann. Math.* **84**, 293–344 (1966)
57. Kleiman, S.L.: The transversality of a general translate. *Comp. Math.* **28**, 287–297 (1974)
58. Knutson, A., Rosu, I.: Equivariant K-theory and equivariant cohomology. *Math. Zeit.* **243**, 423–448 (2003)
59. Kung, J.: Strong maps. In: *Theory of Matroids*, pp. 224–253. *Encyclopedia of Mathematics and its Applications*, vol. 26. Cambridge University Press, Cambridge (1986)
60. Lafforgue, L.: *Chirurgie des grassmanniennes*. American Mathematical Society, Providence, RI (2003)
61. Lazarsfeld, R.: Positivity in algebraic geometry. I. *Ergebnisse der Mathematik und ihrer Grenzgebiete* (3), vol. 48. Springer, Berlin (2004)
62. Lee, S.H., Vakil, R.: Mnëv-Sturmfels universality for schemes. In: *A Celebration of Algebraic Geometry*, pp. 457–468. *Clay Mathematics Proceedings*, vol. 18. American Mathematical Society, Providence, RI (2013)
63. Len, Y.: A note on algebraic rank, matroids, and metrized complexes. *Math. Res. Lett.* (2014). Preprint: [arXiv:1410.8156](https://arxiv.org/abs/1410.8156) (to appear)

64. Lenz, M.: The f -vector of a realizable matroid complex is strictly log-concave. *Adv. Appl. Math.* **51**, 543–545 (2013)
65. Lorscheid, O.: A blueprinted view on \mathbb{F}_1 geometry, chapter. In: *ECM Monograph Absolute Arithmetic and \mathbb{F}_1 -Geometry* (2013, to appear)
66. Mayhew, D., Royle, G.F.: Matroids with nine elements. *J. Comb. Theory Ser. B* **98**, 415–431 (2008)
67. Mayhew, D., Newman, M., Whittle, G.: On excluded minors for real-representability. *J. Comb. Theory Ser. B* **99**, 685–689 (2009)
68. Mayhew, D., Newman, M., Welsh, D., Whittle, G.: On the asymptotic proportion of connected matroids. *Eur. J. Comb.* **32**, 882–890 (2011)
69. Mayhew, D., Newman, M., Whittle, G.: Is the missing axiom of matroid theory lost forever? *Q. J. Math.* **65**, 1397–1415 (2014)
70. Mayhew, D., Newman, M., Whittle, G.: Yes, the missing axiom of matroid theory is lost forever (2014). Preprint, [arXiv:1412.8399](https://arxiv.org/abs/1412.8399)
71. Mikhalkin, G.: What are the tropical counterparts of algebraic varieties. *Oberwolfach Reports* **5**, 1460–1462 (2008)
72. Orlik, P., Solomon, L.: Combinatorics and topology of complements of hyperplanes. *Invent. Math.* **56**, 167–189 (1980)
73. Orlik, P., Terao, H.: *Arrangements of Hyperplanes*. Springer, Berlin (1992)
74. Oxley, J.: *Matroid Theory*. Oxford Graduate Texts in Mathematics, vol. 21. Oxford University Press, Oxford (2011)
75. Pendavingh, R.A., van Zwam, S.H.M.: Lifts of matroid representations over partial fields. *J. Comb. Theory Ser. B* **100**, 36–67 (2010)
76. Peters, C., Steenbrink, J.: *Mixed Hodge Structures*. Springer, Berlin (2008)
77. Poonen, B.: Hilbert’s Tenth Problem over rings of number-theoretic interest (2003). Available at <http://www-math.mit.edu/~poonen/>
78. Proudfoot, N., Speyer, D.: A broken circuit ring. *Beitr. Algebra Geom.* **47**, 161–166 (2006)
79. Read, R.C.: An introduction to chromatic polynomials. *J. Comb. Theory* **4**, 52–71 (1968)
80. Reiner, V.: *Lectures on Matroids and Oriented Matroids*. 2005. Available at <http://www.math.umn.edu/~reiner/Talks/Vienna05/Lectures.pdf>
81. Richter-Gebert, J.: Mnëv’s universality theorem revisited. *Sém. Lothar. Combin.* **34**, 211–225 (1995)
82. Rota, G.-C.: On the foundations of combinatorial theory. I. Theory of Möbius functions. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete* **2**, 340–368 (1964)
83. Rota, G.-C.: Combinatorial theory, old and new. In: *Actes du Congrès International des Mathématiciens (Nice, 1970)*, vol. 3, pp. 229–233. Gauthier-Villars, Paris (1971)
84. Sanyal, R., Sturmfels, B., Vinzant, C.: The entropic discriminant. *Adv. Math.* **244**, 678–707 (2013)
85. Seymour, P.D.: Matroid representation over $\text{GF}(3)$. *J. Comb. Theory. Ser. B* **22**, 159–173 (1977)
86. Seymour, P.D.: On the points-lines-planes conjecture. *J. Comb. Theory Ser. B* **33**, 17–26 (1982)
87. Shaw, K.M.: A tropical intersection product in matroidal fans. *SIAM J. Discrete Math.* **27**, 459–491 (2013)
88. Shor, P.: Stretchability of pseudolines is NP-hard. In: *Applied Geometry and Discrete Mathematics – The Victor Klee Festschrift*. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 4, 531–554. American Mathematical Society, Providence, RI (1991)
89. Speyer, D.: Tropical linear spaces. *SIAM J. Discrete Math.* **22**, 1527–1558 (2008)
90. Speyer, D.: A matroid invariant via the K-theory of the Grassmannian. *Adv. Math.* **221**, 882–913 (2009)
91. Stanley, R.: *Enumerative Combinatorics*, vol. 1. Cambridge Studies in Advanced Mathematics, vol. 49. Cambridge University Press, Cambridge (1997)

92. Stanley, R.: An introduction to hyperplane arrangements. In: *Geometric Combinatorics*, pp. 389–496. IAS/Park City Mathematical Series, vol. 13. American Mathematical Society, Providence, RI (2007)
93. Sturmfels, B.: On the decidability of Diophantine problems in combinatorial geometry. *Bull. Am. Math. Soc.* **17**, 121–124 (1987)
94. Sturmfels, B.: *Solving Systems of Polynomial Equations*. CBMS Regional Conference Series in Mathematics, vol. 97. American Mathematical Society, Providence, RI (2002)
95. Sturmfels, B., Tevelev, J.: Elimination theory for tropical varieties. *Math. Res. Lett.* **15**, 543–562 (2008)
96. Tutte, W.: A ring in graph theory. *Proc. Camb. Philos. Soc.* **43**, 26–40 (1947)
97. Tutte, W.: A homotopy theorem for matroids, I and II. *Trans. Am. Math. Soc.* **88**, 144–174 (1958)
98. Tutte, W.: All the king’s horses – a guide to reconstruction. *Graph Theory and Related Topics*, pp. 15–33. Academic Press, New York (1979)
99. Vakil, R.: Murphy’s law in algebraic geometry: badly-behaved deformation spaces. *Invent. Math.* **164**, 569–590 (2006)
100. Vámos, P.: The missing axiom of matroid theory is lost forever. *J. Lond. Math. Soc. (2)* **18**, 403–408 (1978)
101. Vezzosi, G., Vistoli, A.: Higher algebraic K-theory for actions of diagonalizable groups. *Invent. Math.* **153**, 1–44 (2003)
102. Welsh, D.J.A.: *Matroid Theory*. Academic Press, London, New York (1976)
103. White, N.L.: The bracket ring of a combinatorial geometry I. *Trans. Am. Math. Soc.* **202**, 79–95 (1975)
104. White, N.L.: The basis monomial ring of a matroid. *Adv. Math.* **24**, 292–297 (1977)
105. White, N.L.: Theory of matroids. In: *Encyclopedia of Mathematics and Its Applications*, vol. 26. Cambridge University Press, Cambridge (1986)
106. White, N.L.: Combinatorial geometries. In: *Encyclopedia of Mathematics and Its Applications*, vol. 29. Cambridge University Press, Cambridge (1987)
107. White, N.L.: Matroid applications. In: *Encyclopedia of Mathematics and Its Applications*, vol. 40. Cambridge University Press, Cambridge (1992)
108. Whitney, H.: On the abstract properties of linear dependence. *Am. J. Math.* **57**, 509–533 (1935)
109. Wilson, R.J.: An introduction to matroid theory. *Am. Math. Mon.* **80**, 500–525 (1973)

Index

A

- Abhyankar valuations, 258
- Adic limit theorem
 - adic tropicalization morphisms, 360–363
 - Berkovich setting, 359–360
- Adic seminorm, 205–206
- Adic tropicalization
 - algebraic Gubler models, 348–350
 - of closed embedding, 351–354
 - formal Gubler models, 350–351
 - Huber analytification, 342–343
 - adic spaces, 341–342
 - adic spectra, 340–342
 - admissible formalmodels, 343–345
 - power bounded sections, 342
 - and metrized complexes, 354–360
 - toric degenerations, combinatorics of, 346–348
 - tropical complexes, 354–355
- Admissible formal models, 343–345
- Affine homomorphism, 13
- Algebraic curves, 117, 313, 333, 365, 376, 385, 389, 400, 402, 403, 414
- Algebraic stacks, 289, 307, 309, 312, 316, 329
- Algebraic varieties
 - analytifications, 11
 - currents on, 22–23
 - differential forms on
 - α bidegree, 17
 - Chambert-Loir and Ducros, 21–22
 - non-Archimedean absolute value, 17
 - tropical chart, 17
 - tropical cycle, 17
 - X^{an} subset of, 18–21
 - moment maps, 11, 13
 - morphism of, 14
 - tropicalizations, 11
- Almost isomorphic envelope, 205
- Almost isomorphisms, 205
- Almost tame extensions, 238, 240
 - almost tame fields, 241–242
 - almost unramified extensions, 238
 - deeply ramified fields, 241
 - defectless case, 239–240
 - discrete case, 239
 - non-discrete case, 238
- Almost tame fields, 241–242
 - deeply ramified fields, 241
 - separable extensions of, 242–243
- Amini’s analogous theorem, 424
- Analytic domains, 125
- Analytic functions, 99
- Analytic moment map, 24
- Analytic spaces, 200
 - generalizations to, 24–29
- Analytification, 100
- Anti-continuous, 176
- Arakelov theory, 1, 427
- Arbitrary analytic space, 16
- Artin fans
 - algebraic applications of
 - Gromov–Witten theory, 313–317
 - logarithmic stable maps, birational invariance for, 317–318
 - logarithmic stable maps, boundedness of, 318
 - Relative Gromov–Witten theory, 313–317

- Artin fans(*cont.*)
 categorical context, 308–309
 definition and basic properties, 307–308
 logarithmic scheme, 309
 non-archimedean geometry of
 analytification of, 329–330
 stack quotients, 331–332
 problem and fix
 failure of functoriality, 310–312
 patch, 312–313
- Artin–Hasse exponentials, 74–78
- Artin–Schreier–Witt isomorphism, 80
- Ascoli’s theorem, 38
- Aubin–Mabuchi energy, 38
- B**
- Baker–Norine Riemann–Roch theorem, 379
- Baker–Norine theorem, 380
- Barycentric subdivision, 303
- Base changes
 compatibility with, 255–256
 Jacobians of metric graphs, 255–256
 Kähler seminorms, 224
 Kähler seminorms and, 253–257
- Basis exchange, 443
- Bergman fans
 cryptomorphic definition matroid, 507
 definition of, 499–501
 Minkowski weight, 504–505
 permutohedral variety, 502–504
 $(r_1; r_2)$ -truncation of, 506
 uniform matroid, 502–504
- Berkovich analytic spaces, 2, 11, 399
 motivation, 193–196
 paper and main results, 198–200
 points and stalks of seminorms
 conservative families, 216–217
 points of sites, 215
 P -seminorms 216
 semicontinuity, 215
 stalks, 215
- real-valued fields
 conventions, 200–201
 K^0 -modules, 205–206, 210–212
 K -vectors spaces, 206–207
 seminormed rings and modules,
 201–205
 unit balls, 208–210
- seminormed algebras vs. banach algebras,
 197
- seminormed sheaves
 operations on, 213–214
 pullbacks, 214
 pushforwards, 214
 quasi-norms, 212
 sheafification, 213
- sheaves, 197–198
 analytic points, 217
 analytic seminorms, 217–218
 coherent sheaves, 219–220
 invertible sheaves, 219
 operations, 221
 O_{XG} local rings of, 218
 pullbacks, 219
 seminormed O_{XG} -modules, 219
 unit balls, 197
- Berkovich analytic theory
 Berkovich curves and skeleta, 391–392
 Berkovich spaces, 389–391
 curves, functions on, 392–394
 divisors, tropicalization of, 392–394
 skeletons, higher-dimensional Berkovich
 spaces, 394–397
- Berkovich curves, PL structures on, 53–54
- Berkovich spaces, 179–183
 birational points, 175–176
 deformation retraction, 183–186
 divisorial and monomial points, 177–179
 models, 176–177
- Berkovich spaces
 definability, application of, 124–130
 homotopy type of, 116–124
 Hrushovski and Loeser’s fundamental
 construction, 108–116
 valued fields, general conventions, 102–108
- Berkovich spectrum, 148
- Bieri–Groves theorem, 12, 24
- Birational points, 175–176
- Bounded categories, 204–205
- Brill–Noether theorem, 405–407
- C**
- Calabi–Yau theorem, 34
- Calabi–Yau variety, 188
- Canonical extension, 109
- Cartesian spaces, 206–207
- Chabauty’s theorem, 416
- Chambert-Loir and Ducros, 1, 2
- Characteristic monoid, 296
- Characteristic polynomial
 intersection theory computations, 509–512
 log-concavity, 507–509
- Chip-firing move, 369
- Ciliberto–Harris–Miranda result, 366
- Clifford’s theorem, 380, 381

Cohen's structure theorem, 178
 Cokernels, 202
 Coleman's bound, Katz–Zureick–Brown refinement, 416–419
 Coleman's theorem, 416
 Complex Monge–Ampère equation, 34–35
 Connectedness Theorem, 187
 Conservative families, 216–217
 Convergence polygons, 52, 54–56
 derivatives of, 61–63
 general curves, 59–61
 Convex hull, 119
 Convexity, 86–89
 Costello's comparison theorem, 317
 Cryptomorphism, 442

D

Decomposition theorems, 52
 Deformation retraction, 183–186
 Deligne–Malgrange formula, 66
 Deligne–Rapoport model, 422
 Deligne's theorem, 277
 Dhar's burning algorithm, 383–384
 Differential equations
 first-order, 58, 63, 83
 p -adic, 52, 55
 properties of, 51
 Dini's theorem, 38
 dlt-models, 192–193
 Dual complexes, 34, 41, 46, 133, 135
 degenerations, 138–141
 theorems, 141–144
 tropical complexes, 135–138
 tropical surfaces, 135–138
 Dwork transfer theorem, 61, 63, 86

E

Enumerative invariants, 313
 Étale logarithmic schemes, 328–329
 Exploded fibrations, 337
 Exploded tropicalization, 339
 Exterior and symmetric powers, 203

F

Faithful tropicalization, 146, 147
 analytic space, tropicalization of, 168–170
 for skeleton, 168
 Faltings' theorem, 415
 Filtered colimits, 203
 Finite energy, 39
 Finite morphisms, ramification of, 71–74

Finite separation theorem, 128
 Formal disc, automorphisms of, 82–86
 Free extension, 456
 Fundamental Theorem of Tropical Geometry, 338

G

Gabber's theory, 235
 GAGA-functor, 148
 Gauss valuations
 characterization of, 246–247
 generalized, 245–246
 Generic fiber, 100
 Generic radius of convergence, 51
 Gieseker–Petri theorems, 408–411
 Gödel's completeness theorem, 277
 Gödel's theorem, 277
 Grassmannians
 basis exchange, 479–480
 K -theoretic matroid invariants, 487–490
 matroid representability, 486
 Plücker coordinates, 479–480
 projective toric varieties, 480–481
 realization spaces, 483–485
 representability and semifields, 486–487
 thin schubert cells, 482–483
 torus orbits and matroid polytopes, 481–482
 universality, 486
 Green's Formula, 7
 Gromov–Witten theory, 313–317
 G -topology, 201

H

Hall's theorem, 448
 Hilbert's Theorem, 79
 Hodge index theorem, 135, 441
 Holomorphic coordinates, 3
 Huber adic space, 276
 Huber's adic spaces, 339
 Hurwitz trees, 86

I

Index formula, 52, 63, 66
 Initial degeneration, 11
 Integral affine structures, 159–160

J

Jacobians of finite graphs
 divisors and linear equivalence, 368–370

- Jacobians of finite graphs (*cont.*)
 limit linear series, 370–371
 line bundles degeneration, 367–368
 Néron models, 370–371
- K**
- K-affinoid, 148
 Kähler norm, 225
 Kähler seminorms, 198
 alternative definition, 222–223
 base change, 224
 basic properties of, 223–225
 comparison homomorphism, 233–234
 definition, 221
 discrete valuation case, 235
 and field extensions, 243–245
 filtered colimits, 225
 fundamental sequences, 223
 Gabber’s theory, 235
 and monomial valuations, 245–248
 of real-valued fields
 basic ramification theory notation, 229
 dense extensions, 230–232
 log differentials, 226–228
 tame extensions, 230
 zeroth homology, 229
 universal properties, 222
- Kato cones, 302
 Kato fans
 algebraic spaces, monoidal analogues of
 nodal cubic, 305–306
 Whitney umbrella, 306–307
 and Kato cones, 302
 and logarithmic regularity, 304–305
 points, 302
 resolution of singularities, 303–304
 schemes, monoidal analogues of, 300–302
- Katz–Rabinoff–Zureick–Brown uniformity
 theorems, 419–420
- Katz–Zureick–Brown refinement
 Coleman’s bound, 416–419
 uniformity theorems, 419–420
- Kawamata–Viehweg Vanishing Kollár’s
 Torsion-free Theorem, 191
- k -definable intervals, 106
- Kiehl’s direct image theorem, 28
- Kleiman–Bertini theorem, 508
- Kodaira vanishing theorem, 38
- Kontsevich–Soibelman skeleton, 174
 birational geometry, log discrepancies in,
 187–188
 computation of, 190–191
 definition of, 188
 weight function, 189–190
 work of, 186–187
- Kummer–Artin–Schreier–Witt theory, 79–82
 Kummer isomorphism, 79, 80
- L**
- Linear series. *See also* Berkovich analytic
 theory
 applications
 Brill–Noether theorem, 405–407
 Gieseker–Petri theorem, 408–410
 maximal rank conjecture, 410–411
 break divisors, 386–389
 Dhar’s burning algorithm, 383–384
 Jacobians of finite graphs
 divisors and linear equivalence,
 368–370
 limit linear series, 370–371
 line bundles degeneration, 367–368
 Néron models, 370–371
- Jacobians of metric graphs
 divisors and linear equivalence,
 373–375
 dual graphs under base change,
 371–373
 tropical Abel–Jacobi map, 375–376
 linear systems, 376–378
 rank-determining sets, 384–385
 reduced divisors, 383–384
 Riemann–Roch 379–381
 specialization theorem, 378–379
 tropical independence, 385–386
 Vertex–Weighted graphs, 381–382
- Lines bundles, 32–33
- Lingering lattice paths, 406
- Locally integrable, 23
- Logarithmically regular, 304
- Logarithmically smooth, 299
- Logarithmic growth, 52
- Logarithmic stable maps
 birational invariance for, 317–318
 boundedness of, 318
- Logarithmic structures, 236, 288, 289
 affine toric varieties and cones, 291–292
 Artin fans, 289–290
 analytification of, 290–291
 and Olsson’s stack, 289
 categorical equivalence, 292–293
 charts of, 297–298
 definition, 296–297
 extended complexes, 294
 extended fans, 293
 fans, 292
 functoriality, 295

- invariant opens, 292
 - inverse images, 297
 - Kato fans, 288–289
 - logarithmic differentials, 298
 - logarithmic smoothness, 299–300
 - notation for monoids, 295
 - Olsson’s stack, 289
 - skeletons, 290
 - toric varieties, 288, 291
 - toroidal embeddings, 288, 294
 - without self-intersections, 294
 - tropicalization, 290
 - unobstructed deformations, 289–290
- M**
- Matroid polytopes
 - definition of, 473–476
 - faces of, 476–478
 - grassmannians
 - basis exchange, 479–480
 - K -theoretic matroid invariants, 487–490
 - matroid representability, 486
 - Plücker coordinates, 479–480
 - projective toric varieties, 480–481
 - realization spaces, 483–485
 - representability and semifields, 486–487
 - thin schubert cells, 482–483
 - torus orbits and matroid polytopes, 481–482
 - universality, 486
 - valuative invariants, 477–479
 - Matroids
 - algebraic matroids, 449
 - closure operation, 443
 - as combinatorial abstractions, 439–441
 - definitions of, 442
 - Fano and non-Fano matroids, 446
 - operations on
 - deletion and contraction, 451–452
 - direct sums of, 452–453
 - duality, 453–454
 - extensions, 455–456
 - inequivalent representations, 461
 - ingleton’s criterion, 461–462
 - maps, 457
 - not representable, 462
 - quotients and lifts, 456–457
 - relaxation, 458
 - representability, 458–460
 - polynomial invariants of
 - characteristic and tutte polynomials, 462–465
 - characteristic polynomial, 467–470
 - characteristic polynomial, motivic definition of, 465–467
 - mobius inversion, 467–470
 - reduced characteristic polynomial, 470–471
 - Whitney numbers, log-concavity of, 472–473
 - Metric graphs, lifting problems for divisors
 - chain of loops, 413–414
 - divisors that do not lift, 414–415
 - hyperelliptic curves, 412–413
 - Metrized complexes
 - adic tropicalization
 - general 1-dimensional case, 355–357
 - R -models, type 5" points and inverse systems of, 357–358
 - tropical complexes, 354–355
 - of curves, 400–402
 - divisors from curves, 402–403
 - theory of limit linear series, 403–404
 - Mikhalkin’s correspondence theorem, 333
 - Minimal model program, 193
 - dlt-models, 192–193
 - essential skeleton, 191–192
 - Minkowski weights, 491–499, 504–505
 - Mnev’s theorem, 437, 485, 486
 - Mnev’s universality theorem, 485
 - Model metrics, 32
 - Moduli spaces, 86, 303, 314, 389
 - Brill–Noether rank, 398–400
 - Cartesian diagram of, 317
 - fans and, 333
 - tropical curves, 396–398
 - of tropical curves, 396–398
 - Moment maps, 10–16
 - Monge–Ampère operator, 33
 - Monodromy theorems, 52
 - Monomial points, 258
 - Multiplicative semi-norm, 100
 - Mumford–Neeman theorem, 425
- N**
- Nagell–Lutz theorem, 51, 55
 - Néron models, 370–371
 - Newton polygons, 53
 - Non-Archimedean absolute value, 2
 - Nonarchimedean curves
 - Artin–Hasse exponentials, 74–78
 - Berkovich curves, PL structures on, 53–54

Nonarchimedean curves (*cont.*)

- convergence polygons, 54–56
 - derivatives of, 61–63
 - general curves, 59–61
- finite morphisms, ramification of, 71–74
- formal disc, automorphisms of, 82–86
- index, 63–70
- Kummer–Artin–Schreier–Witt theory, 79–82
- Newton polygons, 53
- subharmonicity, 63–70
- Witt vectors, 74–78

Non-Archimedean dynamics, 1

Non-archimedean fields

- Berkovich spaces, 147–149
- conventions, 147
- curves, case of, 151–152
- faithful tropicalizations
 - analytic space, 168–170
 - for skeleton, 168
- notation, 147
- skeleta, 162–164
 - integral affine structures, 159–160
 - semistable pairs, 160–162
- skeleton, functions on, 164–167
- tropicalization, 149–151

Non-Archimedean geometry, 186, 196

- of Artin fans
 - analytification of, 329–330
 - stack quotients and tropicalization, 331–332

Non-Archimedean Monge–Ampère equations, 34–35, 35

- complex Monge–Ampère equation, 34
- curves, 41–43
- differentiability, 40–41
- energy, 38–40
- envelopes, 40–41
- lines bundles, 32–33
- Monge–Ampère operator, 33
- orthogonality, 40–41
- outlook, 45–47
- singular semipositive metrics, 36–38
- toric varieties, 43–45
- variational approach, 35–36

Non-Archimedean Monge–Ampère measures, 2

Non-Archimedean potential theory, 1, 2

Non-expansive category, 201–202

Non-expansive homomorphisms, 201

Nonzero rational function, 401

O

- Ogg’s theorem, 421, 422
 - $\Omega^{\log} \mathbf{L}^{\circ}/\mathbf{K}^{\circ}$
 - dense extensions, 236–237
 - log different, 237
 - separable extensions, 236
 - tame extensions, 236
- $\Omega_{X/S}$ metrization of
 - base changes, 253–257
 - Kähler seminorms, 248–253, 253–257
- Orthogonal bases, 206–207
- Orthogonality property, 41

P

- p -Adic Cauchy theorem, 55
- Pappus’s theorem, 447
- Permutohedral variety, 441
- Piecewise-linear space, 125
- Piecewise monomiality, 264–265
- PL subspaces
 - integral structure, 261
 - invariants
 - curve, classification of points, 258
 - extensions of valued fields, 257
 - monomial points, 258
 - of points, 257
 - model case, 258–259
 - monomial charts, 259–260
 - skeletons of, 260
 - rational PL-subspaces, 261
 - R_S -PL structures, 260
 - semistable formal models
 - schemes, 261–262
 - skeletons, 262–263
- Plücker coordinates, 400
- Plücker relations, 480
- Pluricanonical forms
 - Kähler seminorms, monomiality of, 263–265
 - maximality locus of
 - residually tame coverings, 266–267
 - residue characteristic zero, 267
 - semistable case, 265–266
 - torus case, 265
 - weight norm, Mustaþa–Nicaise
 - comparison, 269–270
 - weight seminorm, 268
- Plurisubharmonic functions, 33
- Poincaré–Lelong formula, 2

- Polyhedra, 2
- Polyhedral chart, 125
- Polyhedral complex, 7–10
- Polyhedron, 5
- Polytopes, 2

- Q**
- Quotient seminorms, 202

- R**
- Radius of convergence, 51, 55–57, 77, 78
- Realization spaces, 483–484, 484
- Real-valued case, 200
- Relative Gromov–Witten theory, 313–317
- Residually tame morphisms, 255
- Riemann–Hurwitz formulas, 84
- Riemann–Roch theorem, 115, 126, 379–382, 388, 402, 404
- Riesz representation theorem, 23
- Robba’s index formula, 66
- Robertson–Seymour graph minors theorem, 458
- r-orthogonal, 207

- S**
- Schubert varieties, 482
- Seminormed abelian groups, 201
- Seminormed modules, 202
- Seminormed rings, 202
 - completed differentials, 225–226
 - Kähler seminorms
 - alternative definition, 222–223
 - base change, 224
 - basic properties of, 223–225
 - definition, 221
 - filtered colimits, 225
 - fundamental sequences, 223
 - universal properties, 222
- Seminormed sheaves
 - operations on, 213–214
 - pullbacks, 214
 - pushforwards, 214
 - quasi-norms, 212
 - sheafification, 213
- Semistable degeneration, 133
- Semistable pairs, 160–162
- Separation theorem, 130
- Sheaves, 197–198
 - analytic points, 217
 - analytic seminorms, 217–218
 - coherent sheaves, 219–220
 - generization homomorphisms, 275
 - invertible sheaves, 219
 - operations, 221
 - O_{XG} local rings of, 218
 - pullbacks, 219
 - residue fields, 275
 - seminormed O_{XG} -modules, 219
 - spectral seminorm, 274–275
 - stalks, 275–276
 - valuation, 276
- Shilov boundary, 148
- Shokurov–Kollár Connectedness Theorem, 191
- Singular semipositive metrics, 36–38
- Skeleta, 162–164
- Skeletons
 - Berkovich spaces, 318–321
 - faithful tropicalization, 168
 - functions on, 164–167
 - logarithmic structures, 290
 - monomial charts, 260
 - semistable formal models, 262–263
- Skeletons fans
 - and moduli spaces, 333
 - and tropicalization over non-trivially valued fields, 332
- Slope formula, 146, 147, 165–167, 374
- Specialization theorem, 403
- Stack quotients, 331–332
- Stepanov’s theorem, 182
- Stokes’ theorem, 9, 21
- Stone–Weierstrass theorem, 18
- Strictly admissible homomorphisms, 202
- Strictly K-affinoid, 148
- Structure morphism, 296
- Sturmfels–Tevelev multiplicity formula, 14, 26
- Subharmonicity, 63–70
- Supercurrents, 3–7
- Superforms, 3–7
 - on polyhedral complexes, 7–10

- T**
- Tensor products, 203
- Thematic bibliography, 90
- Thin Schubert cell, 441
- t -Monomial valuations, 247–248
- Topological space ${}^1_2X_G/{}^1_2$, 271
 - combinatorial valuations, 278
 - G -skeletons, 277
 - monomiality, 277–278
 - n -dimensional affine R_S -PL space, 277
 - residually unramified monomial charts
 - charts, construction of, 282–283

- Topological space (*cont.*)
 at non-analytic points, 280
 residually unramified locus, 280–281
 unramified extensions, 280
- sheaves
 generalization homomorphisms, 275
 residue fields, 275
 spectral seminorm, 274–275
 stalks, 275–276
 valuation, 276
- structure sheaf, 278
- topological realization
 abundance of points, 271
 germ reduction, 274
 non-analytic points, 273
 notation, 273
 prime filters, 270
 retraction, 272–273
 ultrafilters, 271–272
- torus, G-skeleton of, 279
 valuation monoids, 278–279
 valutive interpretation of points, 279
- Torelli theorem, 376
- Toric degenerations, combinatorics of,
 346–348
- Toric singularities, 300
- Toric varieties, 43–45
 intersection theory, 494–499
 Minkowski weights, 491–494, 494–499
 rational fans, 491–494
- Tropical Abel–Jacobi theorem, 376
- Tropical algebraic geometry, 2, 400
- Tropical charts, 10–16
- Tropical Clifford’s theorem, 380–381
- Tropical cycle, 9
- Tropical geometry, 9, 11, 149, 153, 288, 333,
 337
 fundamental theorem of, 338
- Tropical grassmannians
 dense torus orbit, 155–159
 setting, 152–153
 tropicalization map, 153–155
- Tropicalization
 adic tropicalization
 adic spaces, 341–342
 adic spectra, 340–342
 admissible formal models, 343–345
 algebraic Gubler models, 348–350
 of closed embedding, 351–354
 formal Gubler models, 350–351
 Huber analytification, 342–343
 and metrized complexes, 354–360
 morphisms, 360–363
 power bounded sections, 342
 toric degenerations, combinatorics of,
 346–348
 tropical complexes, 354–355
- algebraic varieties, 11
 artin fans, 331–332
 of Berkovich analytic theory, 392–394
 Berkovich spaces, 318–321
 exploded tropicalization, 339
 faithful tropicalization, 146, 147
 analytic space, tropicalization of,
 168–170
 for skeleton, 168
 logarithmic structures, 290
 map, 24, 359
 toric varieties, 321–324
- Tropical Laplacian, 513
- Tropical linear series. *See* Linear series
- Tropical Riemann–Roch theorem, 409
- Tropical varieties, 145, 146
- Turrittin–Levelt–Hukuhara decomposition, 62
- Tutte’s Homotopy Theorem, 460
- U**
- Unit balls, 197, 204
 almost unit balls, 208
 bounded semilattices, 209–210
 semilattice, 208
 semilattices, 208
- Universally spectral norms, 254
- V**
- Valued fields, 200
- Very affine, 14
- Virtual local index, 69
- W**
- Wagner’s theorem, 458, 459
- Weak Factorization Theorem, 186
- Weierstrass point, 423
- Weierstrass points
 distribution of, 424–427
 on modular curves, 421–422
 specialization of, 423–424
- Weight function, 189–190
- Whitney’s 2-isomorphism theorem, 447
- Whitney umbrella, 306–307
- Witt vectors, 74–78
- Z**
- Zariski logarithmic schemes, 324–328